

U.S. DEPARTMENT OF COMMERCE
National Technical Information Service

AD/A-008 845

A SYSTEM OF INTEGRATED COMPUTER AIDS FOR DECISION
MAKING

RONALD A. HOWARD, ET AL

STANFORD RESEARCH INSTITUTE
MENLO PARK, CALIFORNIA

MARCH 1975

KEEP UP TO DATE

Between the time you ordered this report—which is only one of the hundreds of thousands in the NTIS information collection available to you—and the time you are reading this message, several *new* reports relevant to your interests probably have entered the collection.

Subscribe to the **Weekly Government Abstracts** series that will bring you summaries of new reports as soon as *they are received by NTIS* from the originators of the research. The WGA's are an NTIS weekly newsletter service covering the most recent research findings in 25 areas of industrial, technological, and sociological interest—invaluable information for executives and professionals who must keep up to date.

The executive and professional information service provided by NTIS in the **Weekly Government Abstracts** newsletters will give you thorough and comprehensive coverage of government-conducted or sponsored re-

search activities. And you'll get this important information within two weeks of the time it's released by originating agencies.

WGA newsletters are computer produced and electronically photocomposed to slash the time gap between the release of a report and its availability. You can learn about technical innovations immediately—and use them in the most meaningful and productive ways possible for your organization. Please request NTIS-PR-205/PCW for more information.

The weekly newsletter series will keep you current. But *learn what you have missed in the past* by ordering a computer **NTISearch** of all the research reports in your area of interest, dating as far back as 1964, if you wish. Please request NTIS-PR-186/PCN for more information.

WRITE: Managing Editor
5285 Port Royal Road
Springfield, VA 22161

Keep Up To Date With SRIM

SRIM (Selected Research in Microfiche) provides you with regular, automatic distribution of the complete texts of NTIS research reports *only* in the subject areas you select. SRIM covers almost all Government research reports by subject area and/or the originating Federal or local government agency. You may subscribe by any category or subcategory of our WGA (**Weekly Government Abstracts**) or **Government Reports Announcements and Index** categories, or to the reports issued by a particular agency such as the Department of Defense, Federal Energy Administration, or Environmental Protection Agency. Other options that will give you greater selectivity are available on request.

The cost of SRIM service is only 45¢ domestic (60¢ foreign) for each complete

microfiche report. Your SRIM service begins as soon as your order is received and processed and you will receive biweekly shipments thereafter. If you wish, your service will be backdated to furnish you microfiche of reports issued earlier.

Because of contractual arrangements with several Special Technology Groups, not all NTIS reports are distributed in the SRIM program. You will receive a notice in your microfiche shipments identifying the exceptionally priced reports not available through SRIM.

A deposit account with NTIS is required before this service can be initiated. If you have specific questions concerning this service, please call (703) 451-1558, or write NTIS, attention SRIM Product Manager.

This information product distributed by

NTIS

U.S. DEPARTMENT OF COMMERCE
National Technical Information Service
5285 Port Royal Road
Springfield, Virginia 22161

ADA008845

SEMI-ANNUAL TECHNICAL REPORT
SOFTWARE DEVELOPMENT FOR DECISION ANALYSIS

Contract Number: MDA903-74-C-0240
DARPA Order Number: 2742
Program Code Number: PD420
Contractor: Stanford Research Institute
333 Ravenswood Avenue
Menlo Park, California 94025
Effective Date of Contract: 1 June 1974
Contract Expiration Date: 30 September 1975
Amount of Contract: \$155,000
Principal Investigator: James E. Matheson
(415) 326-6200, ext. 4450
Project Scientist: Allen C. Miller, III
(415) 326-6200, ext. 2865
SRI Project Number: 3309
SRI Proposal Number: MSU-73-133
Contract Period Covered: 1 December 1974 - 1 March 1975

Sponsored by
Defense Advanced Research Projects Agency
DARPA Order Number 2742

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advances Research Projects Agency or the U. S. Government.

Reproduced by
NATIONAL TECHNICAL
INFORMATION SERVICE
U.S. Department of Commerce
Springfield, VA 22151

PRICES SUBJECT TO CHANGE

SEMI-ANNUAL TECHNICAL REPORT
SOFTWARE DEVELOPMENT FOR DECISION ANALYSIS

This report describes research carried out by the SRI Decision Analysis Group for the Defense Advanced Research Projects Agency (DARPA) under contract number MDA903-74-C-0240 during the period 1 December 1974 to 1 March 1975. This research effort has two objectives:

Task A - Develop a morphology for characterizing and analyzing decision problems to serve as a basis for the design of a system of integrated computer aids for decision making.

Task B - Transfer the existing SRI CTREE program to two computers, one available for classified work in the Washington, D. C. area and one accessible through the DARPA computer network.

The attached document describes the research under Task A. A report summary is contained therein.

ia



STANFORD RESEARCH INSTITUTE
Menlo Park, California 94025 · U.S.A.

A SYSTEM OF INTEGRATED COMPUTER
AIDS FOR DECISION MAKING

INTERIM REPORT
DEVELOPMENT OF A DECISION MORPHOLOGY

by

Ronald A. Howard

James E. Matheson

Miley W. Merkhofer

Allen C. Miller, III

Thomas R. Rice

This research was supported by the
Advanced Research Projects Agency
of the Department of Defense under
Contract No. MDA 903-74-C-0240

TABLE OF CONTENTS

	<u>Page</u>
SUMMARY	1
SECTION 1: Research Program	4
1.1 Research Objectives	4
1.2 Research Plan	6
Morphology Development	7
Pilot Experimentation	7
Language Design	9
System Design and Integration	9
Prototype Implementation	9
Test and Evaluation	9
Design for Utilization	10
Communication with Potential Users	10
1.3 Characterizing Decision Problems and Processes .	11
Types of Decision Problems	11
Decision Processes	13
SECTION 2: Components of a Decision Morphology	15
Inferential Notation	16
2.1 Modeling	18
2.1.1 Algebraic and Graphical Languages for Constructing and Describing Models . .	18
Entities and Operators	19
Hierarchical Definition of Operators	20

TABLE OF CONTENTS (Continued)

	<u>Page</u>
Specifying Linkages Between Entities and Operators	23
Directional Nature of Operators	27
Computational Graphs	29
2.1.2 Management of Model Growth	33
2.2 Theoretical Properties of Probabilistic Models .	35
2.2.1 Probabilistic Dependence	38
Independence Equation for Two Random Variables	39
Independence Equation for Three Random Variables	40
Relationships Among Independence Equations	42
Combinations of Independence Equations . .	45
An Additional Relationship Among the Independence Equations	57
Encoding Twenty-two Combinations of Independence Equations	58
2.2.2 Coalescence	63
Influence Diagrams	64
Number of Required Probability Assessments	68
Direction of Arrows in Influence Diagrams	71
Generating a Decision Tree from an Influence Diagram	73
2.2.3 The Value of Decision Dependent Clairvoyance	79
The Primary Decision	79
Value of Perfect Information	82
The Concept of Clairvoyance	86
Practical Probability Assignment	91
Further Implications of the Probability Assignment	93
Decision Dependent Clairvoyance	94

TABLE OF CONTENTS (Continued)

	<u>Page</u>
Unequal Decision Dependent Probabilities	97
Conclusion	98
2.2.4 The Value of Sequential Information	100
2.2.5 The Value of Flexibility	105
2.2.6 Topics for Further Research on the Properties of Probabilistic Models	109
2.3 Analysis of Decision Models	110
2.3.1 Sensitivity Analysis	111
2.3.2 Analysis of Decision Trees	113
2.3.3 Approximate Methods of Probabilistic Processing	116
Background	116
Computerizing the Single Stage Model	119
Steps to Compute the Approximate Value of Information	120
Discrete Decision Variable	121
Continuous Decision Variable	121
Two-Stage Problems	124
2.4 Elicitation of Subjective Information	128
2.4.1 Probability Encoding	129
Individual Probability Encoding	131
Encoding Probabilistic Dependencies	134
Encoding Rare Events	136
Group Encoding	137
2.4.2 Encoding Values and Preferences	139
2.4.3 Encoding Time and Risk Preference	141
2.4.4 Generation and Elicitation of Alternatives	144
SECTION 3: An Illustrative Implementation of the Morphology: A Computerized Graphical-Algebraic Technique for Model Building	146

TABLE OF CONTENTS (Continued)

	<u>Page</u>
3.1 Characteristics of a Successful System of Computer Aids for Decision Analysis	146
3.1.1 System Language Flexibility	147
3.1.2 Multi-level Operation	149
3.1.3 Analytical Capabilities for Decision Analysis	151
3.2 A Graphical Language for Model Development . . .	153
3.2.1 CADMUS (Computer Aids for Decision Makers: User System	158
A Brief Description of the Sulfur Problem	159
3.2.2 Using CADMUS to Analyze the Sulfur Problem	161
Representation of the Deterministic Value Model	169
Problem Analysis	172
3.2.3 Conclusions Concerning the Applicability of the Branching Structure to Single Stage Decision Models	183
3.3 Application of the Graphical Language for Model Development to a Sequential Probabilistic Prob- lem with a Discrete Event Structure.	183
3.3.1 Deterministic Analysis	193
3.3.2 Influence Specification	197
Specifying Information Flow with CADMUS .	198
3.3.3 Constraint Statements	202
3.3.4 Probability Specification	202
3.3.5 Representation of the Decision in a Decision Tree	207
3.3.6 Utility Specification	215
3.4 Conclusions and Summary	215
REFERENCES	216

LIST OF FIGURES

<u>Figure</u>		<u>Page</u>
1.1	Development of Decision Aids, Program Plan	8
2.1	Illustration of Operators and Entities Using a Demand Curve	21
2.2	Three Representations of the Same Algebraic Problem Using Operators and Entities	22
2.3	Models of a Coal Burning Power Plant	24
2.4	Computational Graphs	31
2.5	Model Progression in Decision Analysis	34
2.6	General Description of the Modeling Process in Decision Analysis	36
2.7	Combinations of Independence Equations	46
2.8	Definitions Used in Influence Diagrams	65
2.9	An Example of an Influence Diagram	67
2.10	An Example of a Coalesced Decision Tree	69
2.11	An Application of an Influence Diagram	70
2.12	Another Example of an Influence Diagram	72
2.13	Determining the Order in Which Random Variables and Decisions are Placed in a Decision Tree . . .	76
2.14	An Influence Diagram that Cannot be Translated Directly into a Decision Tree	77
2.15	The Primary Decision Tree	81
2.16	Typical Decision Tree for Determining the Value of Perfect Information	83
2.17	More Complex Tree for Determining the Value of Perfect Information	85
2.18	Clairvoyant's Report Form	87
2.19	Joint Distribution for Clairvoyant's Answers . . .	89
2.20	Decision Tree for Determining the Value of Complete Clairvoyance	90
2.21	Probability Assignment Model	92
2.22	Value of Clairvoyance on Launch Systems Only . . .	95
2.23	Value of Clairvoyance on Mars Only	96
2.24	Sequential Information Decision Tree	102

LIST OF FIGURES (Continued)

<u>Figure</u>		<u>Page</u>
2.25	Decision Regions for Sequential Information . . .	104
2.26	A Portion of a Decision Structure	107
2.27	Classification of Decision Problems	117
2.28	Open and Closed Loop Sensitivities	122
2.29	Compensation for a Discrete Decision Variable . .	123
2.30	Computation for Continuous Decision Variables . .	125
2.31	The Two-stage Quadratic Decision Problem	126
2.32	An Example of Joint Time Risk	143
3.1	About One-third of the Complete Branching Structure for the Sulfur Pollution Problem . . .	156
3.2	Control Panel for the Model Formulation	162
3.3	Step One for Modeling the Sulfur Pollution Problem	164
3.4	Step Two for Modeling the Sulfur Emission Problem	167
3.5	Step Three for Modeling the Sulfur Pollution Problem	168
3.6	Step Four for Modeling the Sulfur Pollution Problem	170
3.7	Step Ten for Modeling the Sulfur Pollution Problem	171
3.8	Sensitivity of Total Societal Cost/KWH to Pollution Cost/lb Sulfur Emitted (High Sulfur Coal)	175
3.9	Sensitivity of Total Societal Cost/KWH to Pollution Cost/lb Sulfur Emitted (Low Sulfur Coal)	177
3.10	Sensitivity of Total Societal Cost/KWH to Pollution Cost/lb Sulfur Emitted (Flue Gas Desulfurization)	178
3.11	Sensitivity of Total Societal Cost/KWH to Pollution Cost/lb Sulfur Emitted	179
3.12	Probability Distribution Illustrating Uncertainty in Health Costs per Pound of Sulfur Emissions . .	181

LIST OF FIGURES (Continued)

<u>Figure</u>		<u>Page</u>
3.13	Probability Distribution Illustrating Uncertainty in Total Annual Societal Costs from a 620 MW Coal Burning Power Plant	182
3.14	Step One for Modeling the Used Car Buyer Problem	186
3.15	Step Two for Modeling the Used Car Buyer Problem	188
3.16	Step Three for Modeling the Used Car Buyer Problem	190
3.17	Step Four for Modeling the Used Car Buyer Problem	192
3.18	Complete Value Structure for the Used Car Buyer Problem	194
3.19	Control Panel for the Deterministic Analysis . .	196
3.20	Used Car Buyer Value Structure Plus Informational Outcomes	200
3.21	Used Car Buyer Value Structure with Appended Influence Diagram	203
3.22	Calculator for the Constraint Specification . . .	204
3.23	Natural Conditioning Order for Probability Specifications for the Used Car Buyer Problem . .	206
3.24	Calculator for Specifying Probabilities	208
3.25	Probability Specification for the Used Car Buyer	209
3.26	Tree Illustrating Joe's Probability Specifications	211
3.27	Decision Tree for the Used Car Buyer Problem	212

LIST OF TABLES

<u>Table</u>	<u>Page</u>
2.1 Summary of Space Mission Example Numerical Results with Original (Equal) and New (Unequal) Probabilities . .	99
3.1 Hierarchy of Automated Decision Aids	150
3.2 Capabilities of Modeling Aid	152

SUMMARY

Decision analysis concepts and methodology have been demonstrated to have high potential for assisting defense decision-makers. However, such assistance for defense executives can now be provided by only a few dozen professionals who have both the required technical competence and communication skills.

Various computerized aids, such as the SRI computer programs for decision tree evaluation, facilitate the analysis of a decision problem; however, the existing software is useful only in the hands of a well-trained analyst or an exceptional executive. Furthermore, these programs only deal with a portion of the decision analysis process.

We are engaged in the first year of a multi-year program to design a broad system of integrated computer aids for decision-making that will foster the widespread application of decision analysis. This system encompasses all of the model building assessment, and analytical steps in the decision analysis process. In particular, the system augments the thought process of the user to describe the problem in convenient terms and to construct a useful decision model. In addition, the system permits assessing the information and preferences of the decision-maker or his delegates. Finally, the system determines the best course of action within the scope of the formulation and, by providing aids such as sensitivity analyses, gives the decision-maker insight into why that course of action was selected.

Rather than begin with the immediate development of various computer aids, we have reviewed the decision process to determine the communicative, creative, and analytic functions that must be performed. Next, we are representing these functions in a precise morphology that constitutes a logical framework for characterizing and analyzing decision problems. The morphology describes the alternatives, probabilities, model structure, values, time preference, and risk preference in the problem as well as the procedures necessary to elicit this characterization and to evaluate the resulting decision model.

Designing computer aids requires striking a compromise between the needs of the user and the difficulty of implementing the procedure on a computer: Too often, the compromise is at the expense of the user. We believe we will ultimately achieve a better design by ignoring the constraints of existing computer systems while developing the morphology. The objectives of the five-year program are therefore:

- The definition of a precise morphology for characterizing and analyzing decision problems.
- The development of a system of integrated computer aids for decision analysis based on the morphology.
- The implementation of the system and its refinement to the point where it is both easy to use and sufficiently powerful to handle large, complex decision problems.

- The communication of the system's advantages to potential users.

This report summarizes our progress to date on morphology definition. It consists of three sections. Section 1 describes our research program to develop a system of computerized aids for decision analysis. Specific research objectives are defined and an eight element, five year program plan for achieving these objectives is outlined. The first section concludes with a discussion of the types and characteristics of decision problems that are encountered in practice. Section 2 of the report deals with the fundamental components of a decision morphology, the analytical and communicative processes that are needed to analyze the broad range of decision problems. Both recent results and areas where further research is essential to the construction and analysis of decision models are addressed here. Section 3 presents an illustrative implementation of the morphology. The graphical algebraic technique described is the precursor to an automated system for decision making and is envisioned as one of its fundamental components.

Our research has exposed many issues material to decision-making, to decision analysis, and to the design of systems for using decision analysis in decision-making. The primary goal of the research during the next six months is to complete the specification of the morphology to the point where it can provide a secure foundation for the development and implementation of a system of integrated computer decision aids.

SECTION 1

RESEARCH PROGRAM

1.1 Research Objectives

The advantage to decision makers of having guidance in the analysis of decision problems has been amply demonstrated by cases where the time and resources for analysis were readily available. However, the need for guidance in rapidly developing decision situations and the shortage of trained analytical personnel mean that few decisions can receive the amount of analysis they deserve.

The major hope in realizing the benefits of insightful analysis lies in using the computer to assist the analytical process. The assistance would be helpful at all levels, from the original characterization of the problem by the decision maker to the sophisticated manipulation of a decision model by the analyst. If such decision aids were available, decision makers could use some of them alone to achieve the benefits of preliminary analysis in problems that now receive only intuitive treatment. Analysts in turn could provide analysis in depth in situations where they now have the time or resources for only initial consideration of important issues.

The availability of effective aids could have an important effect not only on individual decisions, but on the style of decision-making. In our present world without easily accessible decision aids, decision makers can make decisions intuitively without criticism, but they may be considerably, and unconsciously, limited in the size of the projects they will undertake or in their boldness of

execution, perhaps because they lack systematic means to convince themselves or their colleagues of the wisdom of their decision.

The attractiveness of a system that would simultaneously aid in making individual decisions and improve the style of decision making is so great that one would be tempted to start writing computer programs to implement it as soon as possible. However, experience shows that a rush to programming often yields computer aids that reflect existing programming languages more than the needs of the problem. We feel it is important to emphasize the needs of the customer rather than the capabilities of the computer, for the kind of system we are proposing will develop over many years into a complete and powerful workshop of decision aids. Early compromises will become increasingly difficult to accommodate as time passes.

Consequently we have started on this long-term quest of a decision-aiding system by developing a logical framework for characterizing and analyzing decision problems which we call a morphology. Problems posed within the morphology may even be beyond the capabilities of present computers, but this will only restrict the present implementation of the methodology. As computers increase in capability, the type of problem treatable by the morphology will grow apace in complexity and size.

Developing the morphology is thus the first phase of the research program. When it is developed it will be implemented, to

the extent feasible and economical, with present computers. This implementation will in turn require design, execution, and testing of appropriate computer programs. Thus though our ultimate goal is to provide future decision makers and analysts with unusually powerful assistance, we begin by developing a morphology, a formal description of the decision process.

We view the development of a system of integrated aids for decision analysis as a multi-year program guided by the following objectives:

- ° The definition of a precise morphology for characterizing and analyzing decision problems.
- ° The development of a system of integrated computer aids for decision analysis based on the morphology.
- ° The implementation of the system and its refinement to the point where it is both easy to use and sufficiently powerful to handle large, complex decision problems.
- ° The communication of the system's advantages to potential users.

1.2 Research Plan

To accomplish these objectives we have formulated a five-year program plan consisting of eight primary elements. We are currently in the first year of this program. The program elements are:

Morphology Development

Pilot Experimentation

Language Design
System Design and Integration
Prototype Implementation
Test and Evaluation
Design for Utilization
Communication with Potential Users

The time phasing of these elements appears in Figure 1.1;
they have the following brief descriptions:

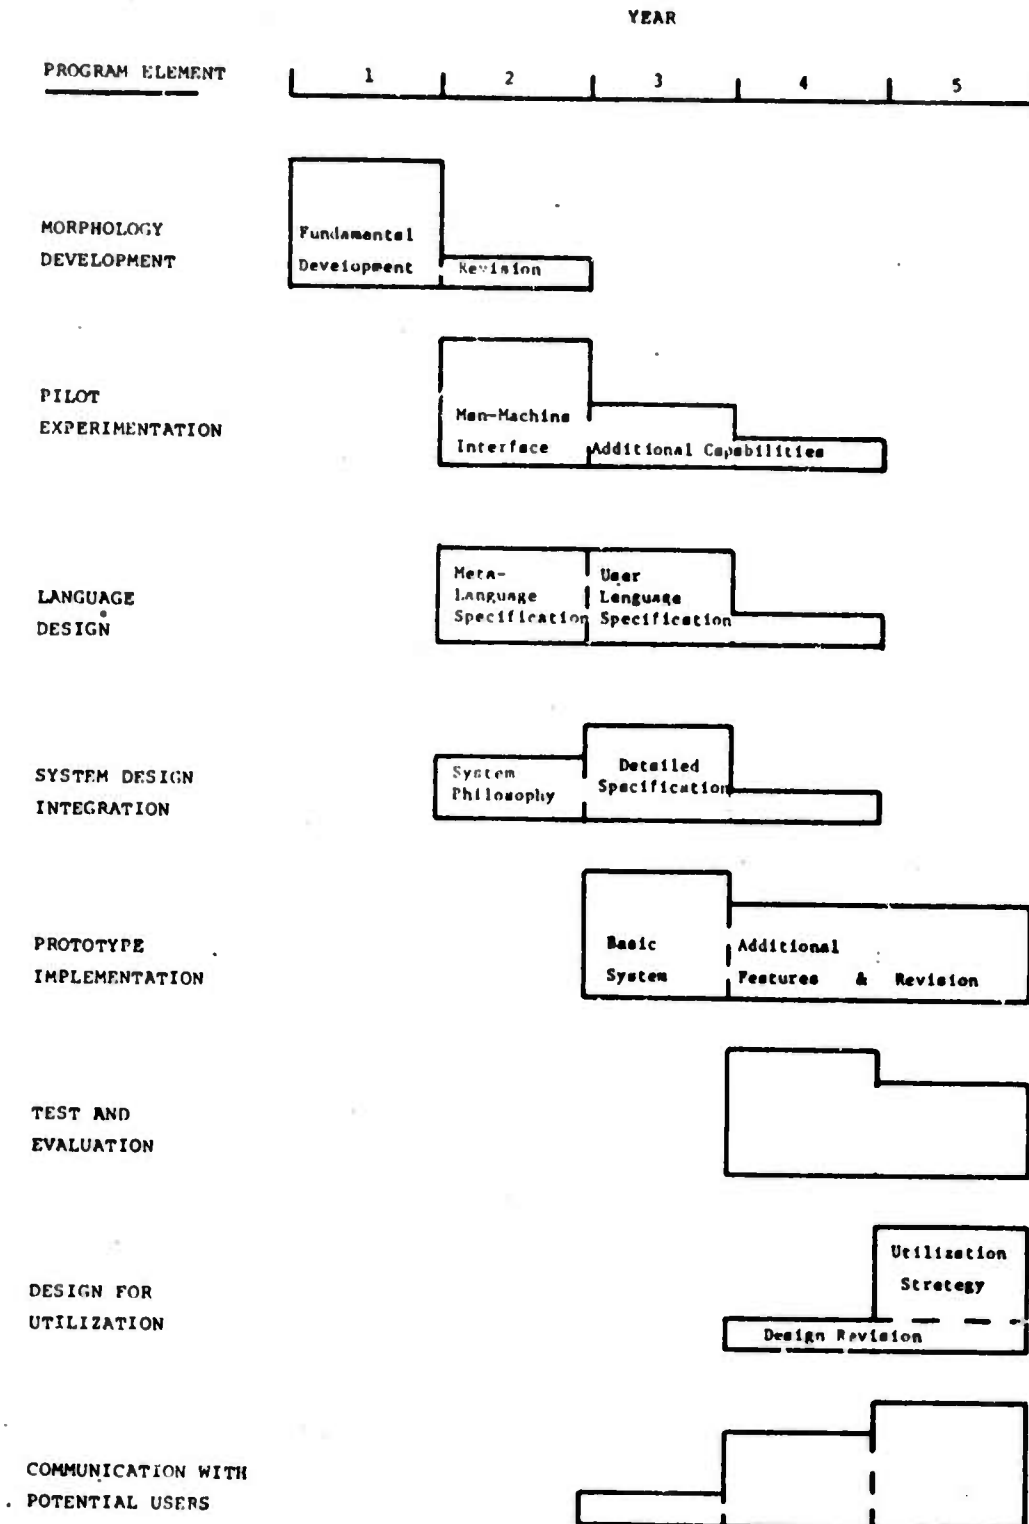
Morphology Development. In this element we review the decision process to determine and describe the communicative, creative, and analytic functions that must be performed. We next define a logical framework for the system of computer aids by developing a precise morphology for characterizing and analyzing decision problems. The morphology must describe the alternatives, probabilities, model structure, values, time preference, and risk preference in the problem as well as define the procedures necessary to elicit this characterization and to evaluate the resulting decision model. The morphology must also treat interactions between each of these components to show how modifications to one part of the process can affect other parts.

Pilot Experimentation. The process of developing the morphology will raise many issues in analysis and elicitation that can

Figure 1.1

DEVELOPMENT OF DECISION AIDS

PROGRAM PLAN



be resolved most easily by experimentation on a limited scale. For example, ease in describing a problem using the morphology can be assessed only in a simulated problem-solving environment. Such experimentation is especially important in choosing means of graphical communication for the user.

Language Design. Once the morphology is specified, it will be possible to design a computer-free meta-language for implementation. This language will then be used to construct the decision-aiding system and the actual language for user-machine communication.

System Design and Integration. The existence of the morphology will permit designing a consistent system configuration and operating philosophy. The design will become increasingly specific in later years through interpretation of the meta-language on the prototype computer system.

Prototype Implementation. In year 3, the prototype system of decision aids will begin to take form on the prototype computer system. The basic system will be ready for testing in year 4. Additional features and revisions will continue throughout the program.

Test and Evaluation. The implemental system will require testing and evaluation in both its basic and advanced forms. Initial tests will be performed on selected system features; later tests will probe the integrating ability of the system.

Design for Utilization. Regardless of its inherent merit, the system of computer aids will achieve its potential only if the process of introducing its use in the decision process is carefully planned. The system must exceed the expectations created for it if user response is to be enthusiastic. Design for utilization will require identifying high priority users and user groups, planning a schedule of introduction consistent with facility and educational requirements, and determining to what extent the prototype system need be modified to meet diverse user needs.

Communication with Potential Users. To create an effective design for utilization will require a multi-year effort to identify users, their needs, and their views on the emergent system of computer decision aids. Joint exploration with potential users should permit selecting a representative group of users for participation in the test and evaluation program.

The five year program will result in a highly effective system of computerized decision aids tested by users and responsive to their needs. We anticipate that this system will be but the beginning of a sequence of progressively more sophisticated decision-aiding systems that will be employed by increasingly diverse and extensive user groups.

1.3 Characterizing Decision Problems and Processes

The preceding discussion has summarized the objectives of this research program and the steps required to achieve the objectives. The following paragraphs deal with the many types of decision problems that are encountered in practice. An immediate objective of our current research is to design a decision morphology that is sufficiently general that it will be relevant to any of the decision problems described below.

Types of Decision Problems. There are many ways to characterize decision problems, but some characteristics are more fundamental than others. During our research we have found the following attributes to be the most useful in describing a decision environment. Decision problems can be characterized by the degree to which each attribute applies.

(1) Time Dependence (static vs dynamic environments):

Decision environments that change over time are inherently more difficult to analyze than those that are static. A situation can be dynamic because parameters are changing (e.g. construction costs are increasing) -- or because the structure of the problem is altered (e.g. a new construction process becomes available).

(2) Sequential decision structure (sequential vs non-sequential decisions). Time can also play an important role in a decision problem if a series of decisions are made sequentially. Each decision in the sequence may have a different structure and

and be based on a different state of information. However, there are often repetitive elements of sequential problems that can be used to simplify the analysis of such problems.

(3) Time availability (planning vs crisis environments):

The time available for analyzing and making a decision often governs the way we approach the problem. The distinction between planning and crisis decisions can be complicated in situations when the time available for a decision is uncertain.

(4) Scale of resources (major vs minor decisions):

There are two resource levels that must be considered in a decision problem: the value of outcomes that might be affected by the decision, and the level of resources allocated by the decision. Decisions can become critical when they involve large amounts of resources, but little time is available for the decision.

(5) Scope of the decision (specific vs policy decisions):

Some decisions deal with specific resource allocations. Other decisions are designed to set policy and give guidance to those who must make specific decisions.

(6) Uncertainty (probabilistic vs deterministic environments):

Decision makers almost never know the exact consequences of choosing an alternative at the time they must make a decision. However, the parameters and structure of a decision problem can be known with varying degrees of expertise.

(7) Complexity (simple vs complex decisions):

Decisions environments can range from those with a few basic elements to those comprised of large numbers of interrelated components. To a certain extent, the complexity of a decision depends on how the decision maker approaches it; one person may take an elementary view of a decision situation while another person could consider the same decision in minute detail.

(8) Continuity (continuous vs discrete environments):

In some decisions, there are only a few discrete alternatives available. In other situations the decision maker can set the level of some parameter to any value in a specified range. Similarly, various elements of the problem can take on only certain values or vary over a range.

Decision Processes. In addition to characterizing decision problems, we can describe the process by which a decision is reached. This process depends primarily on the number of people involved in the decision and their level of experience or training in analyzing decisions.

(1) Number of people involved in decision (individual vs multi-person decision making):

A decision can be made by a single individual, or several people can cooperate or compete to produce the decision. Similarly information and value assessments can be made by individuals or groups. The problem can also be defined or analyzed by

one or more people. In each case, the number of people involved will affect the way in which the situation is analyzed and a decision reached.

(2) Level of training and experience (analytical vs intuitive decision making):

The process by which a decision is reached depends on the backgrounds of the individuals involved in the decision. One person might use highly sophisticated analysis techniques to reach a decision, while another might rely on his intuition.

Similarly those who assess parameters of the decision environment may or may not base their estimates on careful training and encoding procedures.

The following section deals with the analytical and communicative processes that are made to deal with the range of decision problems described above.

SECTION 2

COMPONENTS OF A DECISION MORPHOLOGY

A study of the form and structure of decision analysis is complicated by the fact that there is no single procedure that decision makers and decision analysts follow when dealing with a complicated decision problem. Rather, there are a number of conceptual tools that are available to an analyst, and, within certain limits, the tools can be used in any order.

The conceptual tools of decision analysis are often used iteratively and repeatedly to build an explicit representation of a decision environment. To be effective, a system of automated decision aids should capture this pattern. The decision maker or analyst should be able to move back and forth among the various individual aids selecting the ones that are most useful at each stage of the analysis. For the least experienced users, the order in which the decision aids are employed may be determined automatically on the basis of the user's problem structure. However, more sophisticated users will undoubtedly want a kit from which they can select or construct the most appropriate tool at each stage of the analysis. The purpose of this section is to describe conceptual tools that will form the basis of an integrated set of automated decision aids.

A complete workshop has a wide variety of tools: some are

simple and can be used by the novice, some are highly specialized, and some can be used to build other tools. The same should be true of a comprehensive set of automated decision aids. However, in the following discussion of the conceptual tools that constitute the decision morphology, we will be primarily interested in the more powerful and general aspects of decision analysis that would be of use to an experienced analyst. We must have a clear understanding of these fundamental building blocks before we can develop concepts and decision aids that would be of use to a decision maker with little training in decision analysis.

Some of the following topics have received more attention than others, because we have concentrated our efforts in the areas that are most critical to the development of a unified morphology and the areas that are relatively unexplored. We will now mention only briefly the areas we have not investigated. During the remainder of our research, we intend to examine these areas to balance our understanding.

Inferential Notation

When studying problems of decision making it is desirable to make as explicit as possible the conditions underlying any probabilistic statement. Our discussion of the components of a decision morphology will occasionally make use of a notation, termed inferential notation, that has been developed for this purpose. The following paragraphs are a brief description of this notation.

The basic concept of inferential notation is that every probability assignment is conditional on some state of information,

which we may describe generically by η . If A is some event, we define $\{A|\eta\}$ to be the probability of A given the state of information η . If x is a random variable, then we define $\{x|\eta\}$ to be the density function of x given η . The distinction between events and random variables is usually apparent from context. By extension, if y is another random variable then $\{x,y|\eta\}$ is the joint density function of x and y , $\{x|y,\eta\}$ is the conditional density function of x given y , etc.

A particularly important state of information is the prior experience brought to the problem, which we denote by ξ . Any probability assignment conditional only on ξ is called a prior assignment. Thus, $\{A|\xi\}$ is the prior probability of the event A ; $\{x|\xi\}$ is the prior probability density of the random variable x . For simplicity we sometimes shorten $\{A|\xi\}$ to $\{A\}$ when the state of information is obvious from context.

The notation extends to moments. We define $\langle x|\eta \rangle$ to be the expectation of the random variable x given the state of information η , computed from

$$\langle x|\eta \rangle = \int_x x \{x|\eta\}$$

where \int_x is a general summation operator. Then $\langle x^n|\eta \rangle$ is the n th moment of x . We denote the variance of x by

$$v_{\langle x|\eta \rangle} = \langle x^2|\eta \rangle - \langle x|\eta \rangle^2 = \langle (x - \langle x|\eta \rangle)^2|\eta \rangle.$$

2.1 MODELING

2.1.1 Algebraic and Graphical Languages for Constructing and Describing Models

To conceptualize and construct a model that adequately describes a decision situation, an analyst often attempts to visualize the decision process in both graphical and algebraic terms. For example, the output of a power generating plant can be described as an algebraic function of the amount of fuel consumed by the plant, the energy content of the fuel, and the plant efficiency. Alternatively, the role of the plant in a complex power generating and distribution system might be described graphically by a flow chart that indicates the various material and power flows throughout the system. Similarly, the probabilistic dependence of international events on U.S. foreign policy decisions can be described algebraically by a set of conditional probability distributions, or graphically by a decision tree with event probabilities dependent upon the alternatives chosen by U.S. policy makers.

Some types of problems are easier to conceptualize in algebraic terms; other problems, in graphical terms. The choice of a particular modeling language depends on both the nature of the problem and the background of the analyst. Some analysts prefer to visualize complex systems in terms of flow charts, tree structures, block diagrams, or graphical functions in several dimensions.

Other analysts might deal with the same problem in terms of systems of equations, alpha-numeric lists of data, or algebraic computer programs (FORTRAN, ALGOL, LISP, etc.)

There is no single language that is "best" for describing a particular decision situation. In fact, there is value in looking at several different descriptions of the same problem. Different modeling languages tend to illuminate different aspects of the problem being studied. An analyst may often switch back and forth between algebraic and graphical representations of a decision environment and may also use several different languages of each type in analyzing a single decision problem.

For example, a decision tree is often used to describe and analyze complex decision situations. However, decision trees can become so large and complex that they cannot be visualized, let alone drawn, in their entirety. As a result, analysts have developed a number of compact graphical notations for summarizing the essential features of a large decision tree. In this case, the analyst may find it valuable to translate among several different graphical and algebraic representations of the problem.

Entities and Operators

All modeling languages have in common certain fundamental building blocks. These elements can be combined to describe many complex decisions. The basic elements of all models are entities and operators. Entities describe the state of the environ-

ment, and operators describe how the state of the environment is modified. Modeling entities can be numbers, arrays, functions over time, strings of alpha-numeric data, algebraic variables, logical variables, decisions, complex variables, etc. Operators transform one entity into another.

For example, Entity E_1 could be the demand for a particular quantity at a certain time. Entity E_2 could be the equilibrium price of that quantity in a competitive market. Operator O_1 could transform the demand for this product into the equilibrium market price. An economist might view this transformation in terms of the demand curve shown in Figure 2.1a, but in more general terms, we can view this process as that of an operator (the demand curve) transforming entity E_1 (demand) into entity E_2 (price) as shown in Figure 2.1b.

In general, operators can transform combinations of entities and operators into other combinations of entities and operators. For example, operator O_1 could take values of one entity called X and transform them into another entity called Y , using the equation $Y = 3X + 2$. A second operator, O_2 , could be one that solves algebraic equations for X . When O_2 is applied to O_1 , a third operator is produced, O_3 . O_3 takes the entity called Y and transforms it into the entity called X using the equation $X = (Y - 2)/3$. These three operators are shown in Figure 2.2.

Hierarchical Definition of Operators

Not only can operators transform entities and other operators, but an operator may consist of a set of entities and operators. For

FIGURE 2.1A DEMAND CURVE

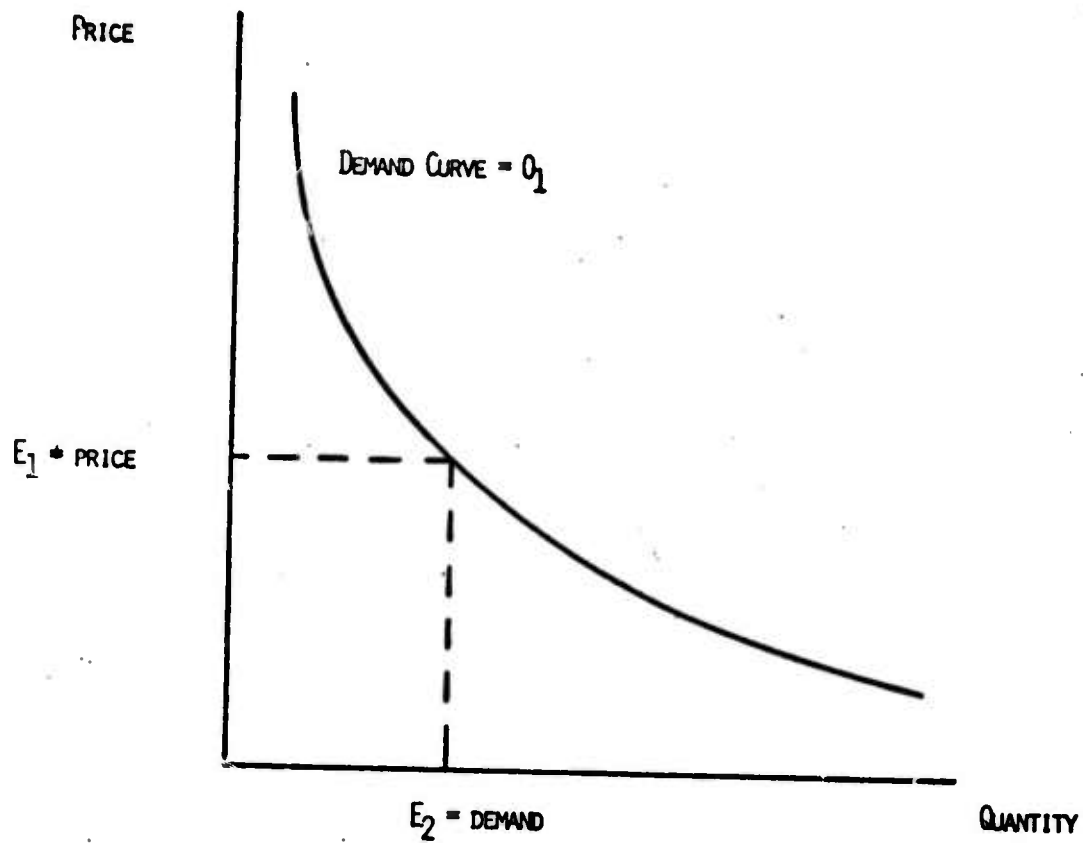


FIGURE 2.1B GENERAL NOTATION

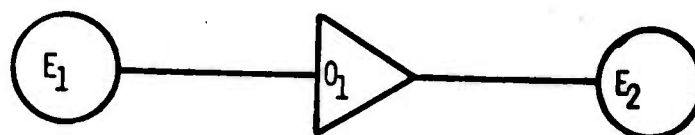


FIGURE 2.1 ILLUSTRATION OF OPERATORS AND ENTITIES USING A DEMAND CURVE

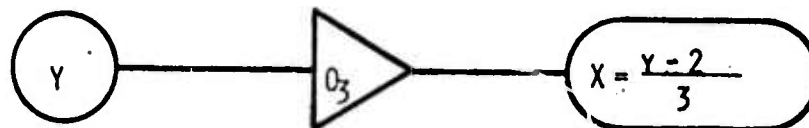
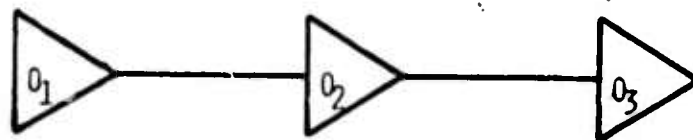
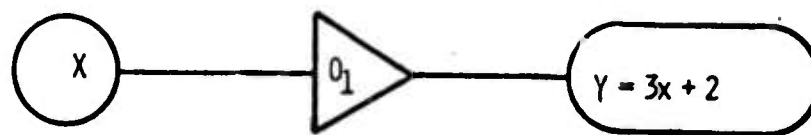


FIGURE 2.2 THREE REPRESENTATIONS OF THE SAME ALGEBRAIC PROBLEM
USING OPERATORS AND ENTITIES

example, a simple model of a coal-burning power plant might consist of a single operator that transforms the fuel consumption of the plant and the energy content of the fuel into the power produced by the plant. The operator that carries out this transformation is shown in Figure 2.3a. However, in a more detailed model of the power plant's performance, this operator may consist of several components that describe the parts of the power plant. Figure 2.3b shows an expanded model in which the operator describing plant's power production has been decomposed into other operators and entities that collectively transform the fuel consumption and energy content of the fuel into the power produced by the plant.

It is often valuable to expand portions of the model by describing the process in greater detail. However, large models are often hard to understand without combining model elements into global units that show the analyst an overview. By allowing operators to comprise other operators and entities, an analyst can aggregate and disaggregate the model and thus obtain perspective on the problem.

Specifying Linkages Between Entities and Operators

The links between operators and entities are often described implicitly by the definition of the operators. For example,

FIGURE 2.3a SIMPLE MODEL

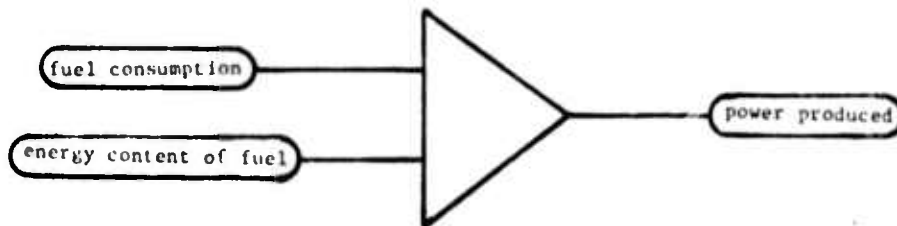


FIGURE 2.3b EXPANDED MODEL

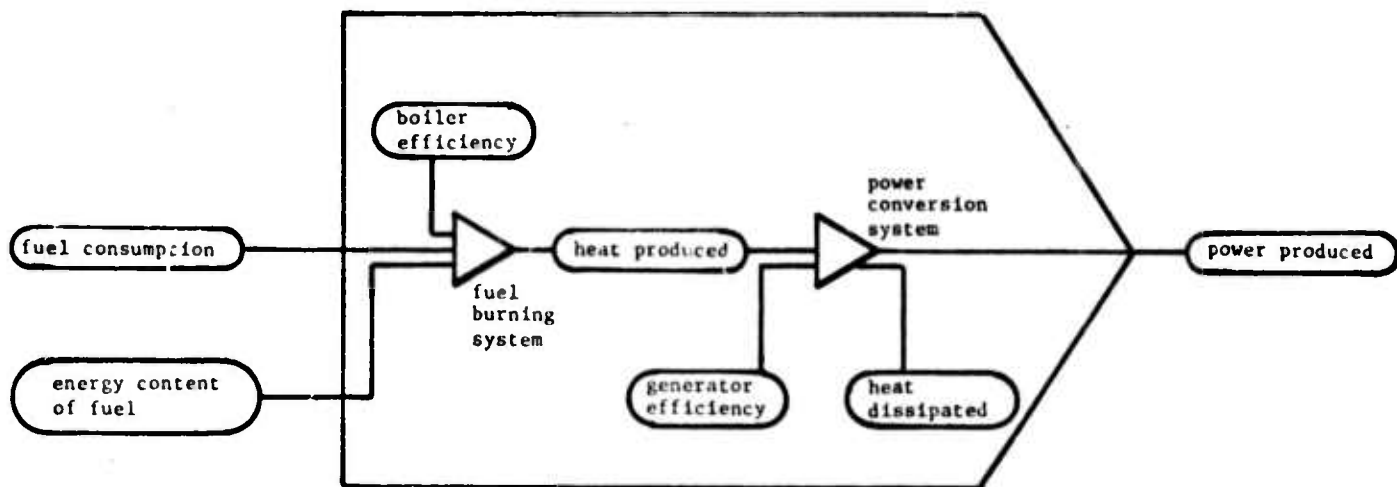


FIGURE 2.3 MODELS OF A COAL BURNING POWER PLANT

an operator that transforms demand into price should be linked to an input entity representing demand and to an output entity representing price. However, the linkages between operators and entities need not be specified by the operator definition. It is possible to find general types of operators and entities, and later specify the linkages that connect them. While a separate specification of the linkages between entities and operators may be rather cumbersome for simple problems, it can be a powerful technique for generating large models.

A typical large decision model might contain thousands of entities and operators. If an analyst were required to specify each of the entities, operators, and linkages, the process of modeling would be long and tedious. Furthermore, the analyst could easily lose sight of the global structure of the problem. To overcome these difficulties, analysts have developed several procedures for specifying large models by specifying classes or hierarchies of entities and operators. For example, combinations of entities and operators are often grouped together to form a composite operator, as shown in Figure 2.3.

However, the definition of operators consisting of other operators and entities is only one of the ways an analyst can define classes or hierarchies of operators. In general, classes of operators can be defined through the use of "operator rules." An operator rule defines a partition of all of the operators and entities that

exist in a particular decision model. For example, an operator rule might state that all of the elements of a particular model are associated with one of several types of materials used in the production of energy (coal, oil, nuclear power, etc.). A second operator rule might state that all entities exist in one of several regions (the Persian Gulf area, North Africa, the United States, Western Europe, etc.). The intersection of these two operator rules defines a large number of entities, such as Persian Gulf oil, Western European coal, etc.

The operator rules may contain a statement that some of the elements defined by the intersection of the two rules cannot exist (for example, Persian Gulf coal). A third operator rule might state, in general terms, the relationship between various classes of entities, for example the relationship might state the amount of gasoline that can be produced from a unit of crude oil. A combination of these rules can be used to define a very large network of entities and operators without making it necessary for the analyst to describe every element and operator in the model.

We can view the generation of a large model through the use of operator rules as a transformation from one modeling language to another. The operator rules are themselves described by general entities and operators that are combined with those of other operator rules to form the complete model. A set of operator rules can, in itself, be a complete model of the decision situation, thus, even at a very basic level, there is no unique model for modeling language.

Directional Nature of Operators

By definition, operators have a direction; they transform inputs into outputs, but not vice versa. Thus, we can view operators as a mapping from one entity space to another. One implication of this definition is that an equation is not an operator unless we specify which variables in the equation are inputs and which are outputs. For example, the equation $X = X + 1$ has no solution, but if we define the left side of the equation to be the output and the right side of the equation to be the input, this equation defines a perfectly valid operator. This operator has entity X as both an input and an output; its function is to add one to the value of X . If we wish to solve algebraic equations using our definition of entities and operators, it is necessary to specify the equation in terms of an operator with an input and an output. We can then define a second operator that solves the equation and produces a third operator with different inputs and outputs (see Figure 2.3).

By defining operators to be directional, we can avoid a number of theoretical difficulties. If we were to view operators as simply being relationships among entities, without the concepts of input and output, we could find ourselves in the position of specifying models that are either insoluble or possess many solutions. Although we can define a relationship in which the entity Y can be determined from the value of entity X , the mapping from X to Y need not be 'one-to-one' or 'bnto'. "One-to-one" means

that there is a unique value of Y for every different value of X . "Onto" means that there is some value of X corresponding to every possible value of Y .

If the mapping from X to Y is not one-to-one and onto, then certain values of Y will have no corresponding values of X . In this case, it is very difficult to interpret the relationship between X and Y as one in which X can be determined from Y . In other words, the relationship between X and Y contains implicit direction: X is the input and Y is the output.

It is possible to reverse the mappings that are not one-to-one and onto by defining certain special conventions, but we believe it would be better to include these conventions in the definition of a directional operator. For example, equation

$$Y = X^2$$

is unambiguous when we are trying to determine Y given X . However, when we are given Y and attempt to solve for X , we find that there can be more than one solution or no solutions at all, depending on the value of Y . In order to overcome this difficulty, we might establish the convention that X will always be the positive square root of Y , if Y is greater than or equal to zero, and X will be zero if Y is less than zero. By making operators directional, we require that these conventions be an explicit part of the operator definition.

If we were to consider operators to be non-directional relationships among entities, then we could no longer consider the

the operator to be a mapping from the space of one entity to the space of another. Instead, operators would define the sub-space of the space spanned by all of the entities associated with the operator. Thus, for example, the relationship $Y = 3X + 2$ defines a straight line in the two-dimensional Euclidean space spanned by X and Y . This simple example might lead us to assume that each operator or relationship defines an $(n-1)$ dimensional sub-space of the n -dimensional space spanned by the entities. However, this is not the case. The equation $X = X$ does not reduce the dimensionality in the space, while the equation $X^2 + Y^2 = 0$ reduces a two-dimensional space to one with zero dimensions. To avoid these conceptual difficulties, we will limit our definition of operators to meaning a directional mapping from a set of inputs to a set of outputs.

Computational Graphs

When operators and entities are linked together, either individually or through the use of operator rules, the result is a computational graph that specifies the structure of the model. These graphs can be used to specify any of the models used in decision analysis: decision trees, Markov processes, financial models, material flows over time, fault trees, etc. In fact, given the very basic nature of operators and entities, it is difficult to conceive of a model that cannot be put into the form of a computational graph. This is not to say that computational graphs are the most efficient way to model every decision problem, but rather that they are sufficiently important that they should be incorporated into any general modeling language.

Computational graphs can be classified according to the degree they are interconnected. One of the simplest forms of computational graphs is a tree. All the linkages of a tree are directed either from or to a unique starting point or origin. There must be only one path that connects any pair of nodes, and if one of those nodes is the origin, all of the operators along that path must be oriented in the same direction. Figure 2.4a shows an example of the tree.

A slightly more general version of a computational graph is a lattice. A lattice is a tree in which the branches are allowed to coalesce or connect together. In a lattice there can be multiple paths between each pair of nodes, but if one of the nodes is the origin, then all of the operators along each of the paths must be oriented in the same direction--either toward the origin or away from the origin. Figure 2.4b shows an example of a lattice.

Trees and lattices have the advantage of not containing loops. A graph contains a loop if two nodes in the graph are connected by two or more paths, such that all of the operators along one path lead from the first node to the second and all of the operators along another path lead in the other direction. An example of a graph containing a loop is shown in Figure 2.4c.

The existence of loops makes it much more difficult to process the computational graph, but it also makes it possible to simplify the structure of the graph. Conventions must be established for dealing with loops in computational graphs, but once this is done,

FIGURE 2.4A - EXAMPLE OF A TREE

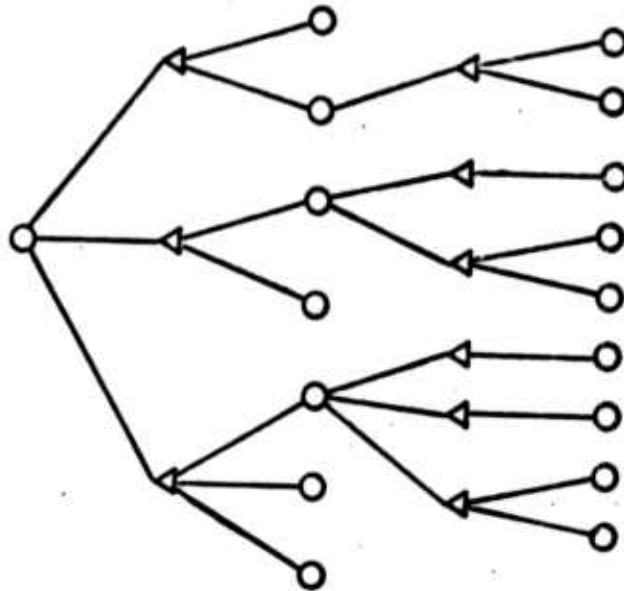


FIGURE 2.4B - EXAMPLE OF A LATTICE

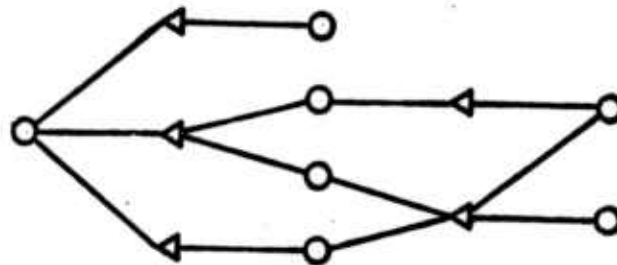


FIGURE 2.4c - EXAMPLE OF A GRAPH WITH A LOOP

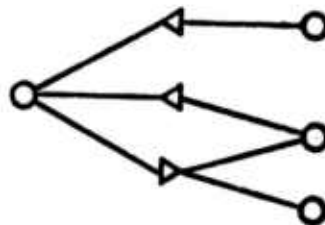


FIGURE 2.4 COMPUTATIONAL GRAPHS

it may be possible to represent very large (in fact, infinite) decision problems in compact form. For example, a Markov process can be represented as an infinitely large decision tree. Although a decision tree may be easier to process than a Markov diagram, the unbounded nature of the decision tree makes it impossible to solve completely. On the other hand, Markov processes can be represented and solved using relatively simple and compact computational graphs containing loops.

These basic components of a modeling language are sufficiently general to allow the construction of many different types of models--both probabilistic and deterministic. An example of how such a language might be implemented is presented in Section 3 of this report.

2.1.2 Management of Model Growth

Historically, modeling has been the art of decision analysis, while decision theory has been the science. Modeling will always require judgment and creativity from the analyst, but there are several areas in which computerized decision aids can be used to help the creative process.

The type of modeling used in decision analysis is "top-down". For example, in government decisions the objective is to maximize social profit; however, directly estimating social profit for each alternative in a decision problem is equivalent to choosing the best alternative by intuition. To effectively use experts we divide social profit into benefits and costs. Each of these categories is further subdivided until we reach the level where experts are available to estimate the inputs to the model. The stopping point is a decision not an outcome. For example, if one of the costs is for crude oil in 1980, we can estimate it directly, or we can build a world energy model to predict the future price of crude oil.

Figure 2.5 shows how top-down modeling is integrated into the decision analysis cycle. At the start of the project intuition and direct assessment are used to predict social profit and many alternatives are discarded without formal analysis. As the deterministic phase progresses increasingly complex models are built. Near the end of the deterministic phase the model complexity peaks. Then deterministic sensitivities are used to eliminate unimportant variables. Finally the model is frozen for use during the probabilistic phase.

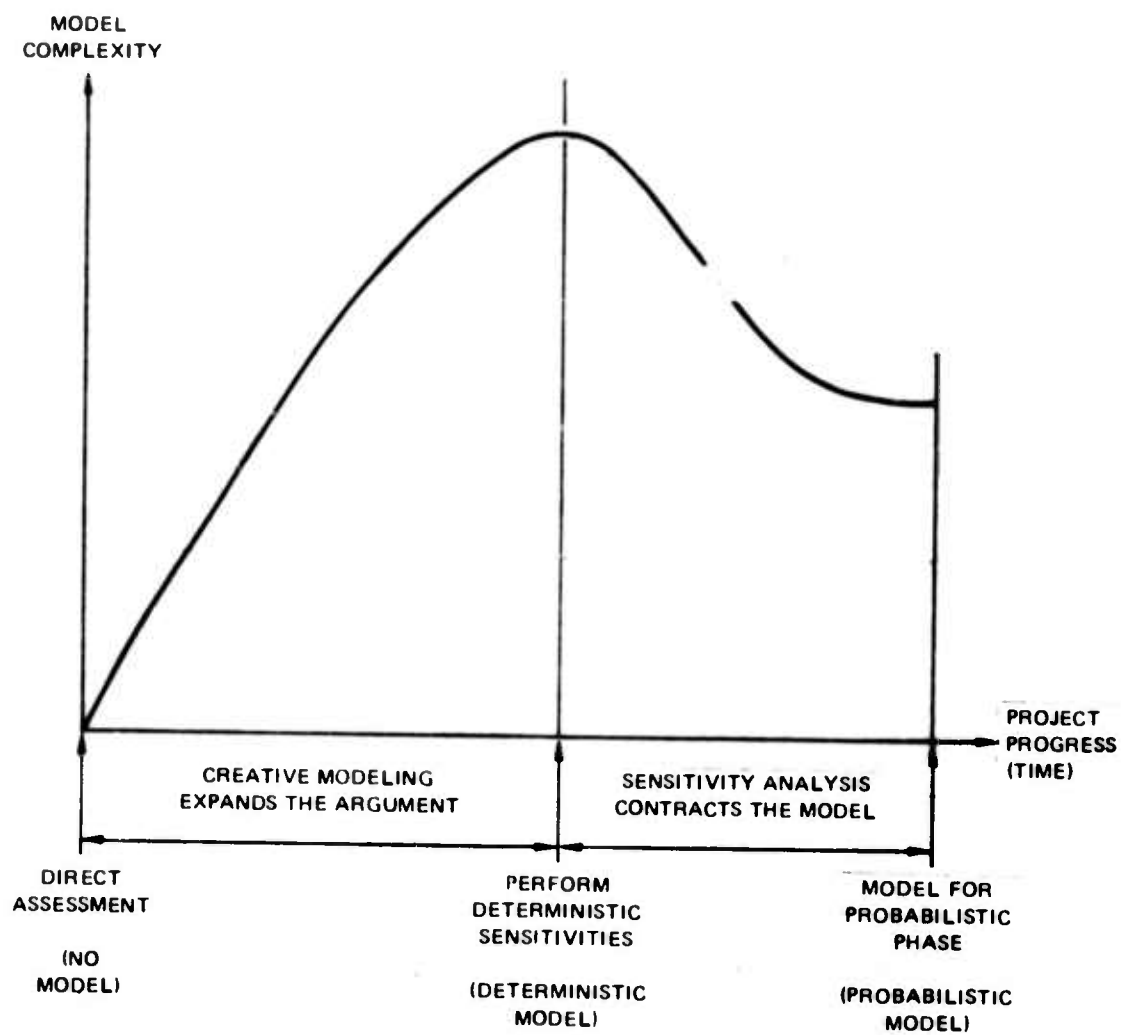


FIGURE 2.5 MODEL PROGRESSION IN DECISION ANALYSIS

A more realistic view of decision analysis is the rising wave of Figure 2.6 rather than the single hump of Figure 2.5. The analyst uses sensitivities both to identify areas where additional modeling is required and to eliminate unnecessary complexity. The result is that while the complexity continually expands and contracts, the average over time gradually rises, capturing only the most important aspects of the decision problem.

Eliminating unnecessary complexity is one distinguishing characteristic of decision analysis. Without our willingness to discard data, model complexity would rise exponentially as shown by the dotted line in Figure 2.6.

The distinction between model expansion and contraction is important because the modeling effort requires different computer aids during each phase. During the construction or expansion of a model, an analyst will use general modeling languages, such as those discussed in the preceding section. To eliminate unnecessary complexity, and limit the size of the model, an analyst needs conceptual tools for evaluating the relative worth of the various components of the model. These tools--like sensitivity analysis, and procedures for determining the value of various types of information and flexibility--will be discussed in the following sections.

2.2 THEORETICAL PROPERTIES OF PROBABILISTIC MODELS

Probabilistic models are fundamental to the analysis of decisions under uncertainty. However, the theoretical development

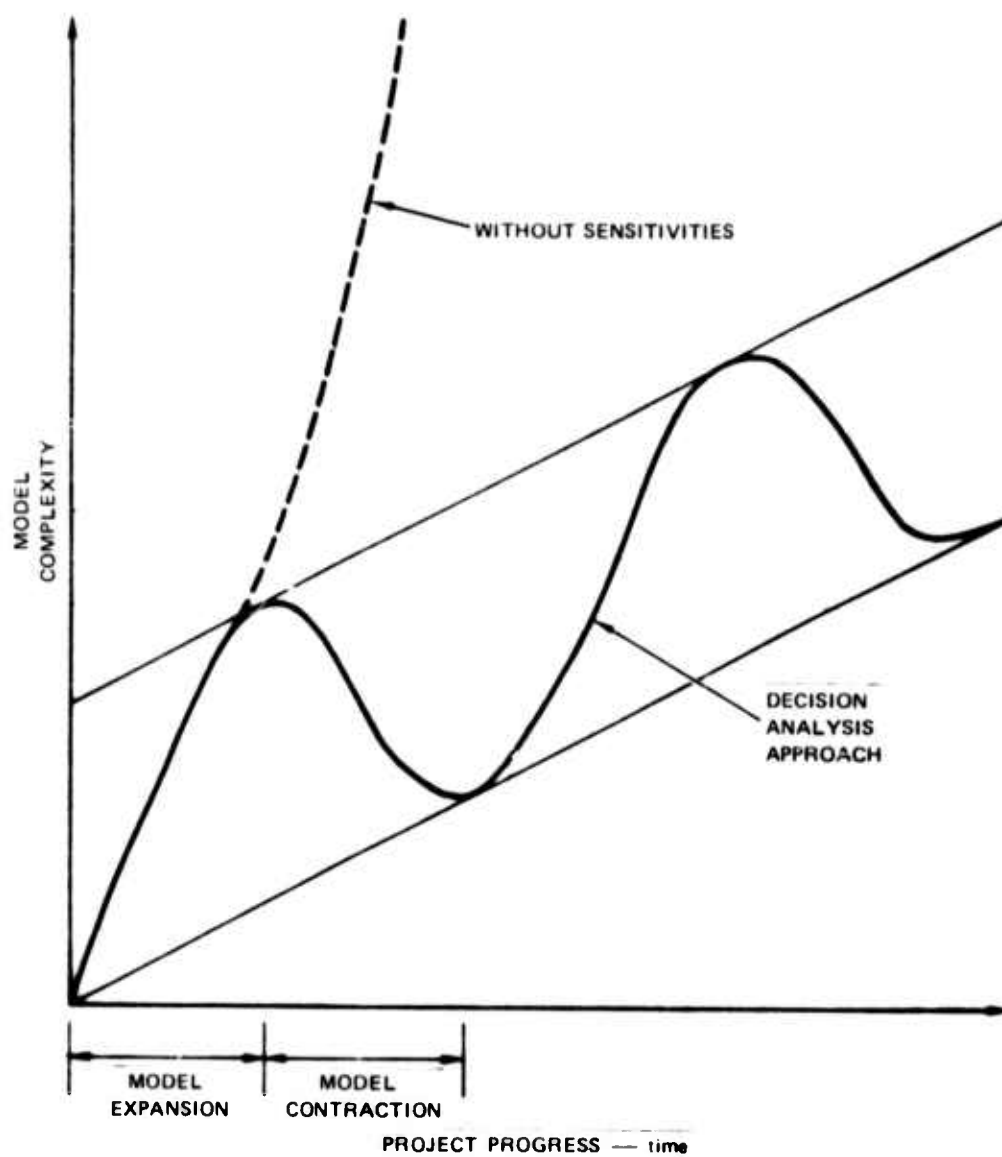


FIGURE 2.6 GENERAL DESCRIPTION OF THE MODELING PROCESS IN DECISION ANALYSIS

of procedures for dealing with probabilistic models has not kept pace with the extensive literature on deterministic models (linear systems theory, optimal control, linear programming, optimization techniques, etc.). The purpose of this section is to explore some of the theoretical properties of probabilistic models that will form the basis for conceptual modeling tools in automated decision aids.

There are two ways to look at probabilistic models. One is to view them as deterministic models that contain some uncertain parameters, in which case the major modeling effort is devoted to building the deterministic model. In this case, most of the relationships among the parameters of the model are in the form of equations; uncertainty is introduced into the model by specifying the probability distributions for various model parameters.

The second approach is to use a probabilistic model as a direct description of an uncertain environment. In this case, the relationships among the parameters of the model are often probabilistic, meaning that even if we knew all of the inputs to the model with certainty the output of the model must still be described by a probability distribution.

It is possible to view the first approach to probabilistic modeling as a special case of the second, in which all of the probabilistic relationships among the parameters of the model are approximated by deterministic equations. In many modeling efforts both approaches are used; the resulting decision models contain

both deterministic and probabilistic parameters, and both deterministic and probabilistic relationships among the parameters.

The following subsections explore the probabilistic relationships that can exist among parameters of a model and the implications of the various uncertainties in the model for the value of information and flexibility. The last subsection contains a brief discussion of two additional theoretical properties of importance to probabilistic modeling: joint time-risk preference, and models of competitive decisions.

2.2.1 Probabilistic Dependence

To simplify the tasks of assessing and processing uncertain information it is often necessary to assume that random variables are independent. With this assumption, it is possible to assess the probability distribution for each random variable separately, and deal with relatively simple marginal probability distributions rather than complicated conditional distributions. Most large decision analysis projects contain within them some implicit or explicit assumptions about the independence of various random variables.

In order to describe independencies, we need to have a clear understanding of the types of independence that are possible. Two random variables are either dependent or independent. However, as we shall see, there are twenty-two different combinations of dependence and independence that can exist among three random variables.

When there are more than three random variables, there are many different combinations of dependence and independence that exist among them.

This section describes the twenty-two combinations of dependence and independence that can exist among three random variables and gives an example of each combination. A similar discussion can be found in a book entitled, "Rational Descriptions, Decisions and Designs," by Myron Tribus [12]. However, Tribus discusses only twelve of the twenty-two possible combinations of dependence and independence that can exist among three random variables.

Independence Equation for Two Random Variables

When we assert that two random variables are independent, we are assuming that their joint probability distribution is equal to the product of the two marginal probability distributions. In other words, if we say that random variables A and B are independent, then we are claiming that the following equation is true.

$$\{AB\} = \{A\} \{B\}$$

Alternatively, if we say that A and B are dependent, then we are claiming that the equation above is not true. To avoid writing an equation for each independence assertion, we will denote the assumption of independence between random variables A and B as follows: $I(A,B)$. If A and B are dependent, then we write $\bar{I}(A,B)$. If there are only two random variables, then we must have either $I(A,B)$ or $\bar{I}(A,B)$.

Independence Equations for Three Random Variables

If we have three random variables--A , B , and C -- then there are ten possible independence equations. For example, we can assume that the three random variables are mutually independent, which means that their joint probability distribution is equal to the product of the three marginal distributions. Mutual independence is equivalent to the following equation.

$$\{ABC\} = \{A\} \{B\} \{C\} .$$

If the three random variables are not mutually independent, then this equation is not true. In our notation mutual independence is written $I(A,B,C)$. If the three random variables are not mutually independent, we write $\bar{I}(A,B,C)$. The lack of mutual independence does not mean that the three random variables are mutually dependent, since other types of independence are possible. Mutual dependence means that there are no independencies among any of the random variables.

Another type of independence that can be asserted among the three random variables is that two of the three random variables are independent without regard to the third. When we have three random variables-- A , B , and C --we can assert $I(A,B)$ or equivalently

$$\{AB\} = \{A\} \{B\} .$$

In our notation, this independence equation is denoted $I(A,BC)$. This equation means that learning the value of A will not change the joint probability distribution for B and C . Alternatively,

it means that learning B and C will not change the probability distribution for A. We could also assert that B is independent of A and C, or that C is independent of A and B.

When we learn the value of one of the three random variables, it changes our state of information and therefore can change our assumptions about independence for the remaining random variables. This allows us to make a different kind of independence assumption. For example, we can assume that when we know C, A and B are independent. The equation for this independence assumption is

$$\{AB|C\} = \{A|C\} \{B|C\}$$

When $\{C\} \neq 0$,* the following equation is equivalent to the one above.

$$\{ABC\} = \frac{\{AC\} \{BC\}}{\{C\}}$$

This independence assumption is denoted $I(A,B|C)$. As before, we can find some other independence assumptions by permuting the random variables. Thus, we can assume that B and C are independent when we know A, or we can assume that A and C are independent when we know B.

In this discussion it will be assumed that independence can only be asserted among quantities that are each based on the same state of information. Independence among quantities that are based on different states of information is difficult to conceive since the different states of information may themselves affect the validity of the independence assumption. For example, can we assert

*This is not a restrictive condition since $\{AB|C\}$, $\{A|C\}$, and $\{B|C\}$ are not defined when $\{C\} = 0$.

independence between A by itself, and B given C ? The answer to this question is no, since the independence equation corresponding to this statement is not well defined. We are trying to assert independence between {A} and {B|C}, and by analogy with the other independence equations we should be able to equate the product of these two quantities to some other quantity. However, it is not clear whether the product should equal {AB}, {AB|C}, or something else.

In summary, there are ten possible independence equations or assumptions that can be made for three random variables. These independence equations are listed below.

- (1) $I(A,B,C) \leftrightarrow \{ABC\} = \{A\} \{B\} \{C\}$
- (2) $I(A,B) \leftrightarrow \{AB\} = \{A\} \{B\}$
- (3) $I(A,C) \leftrightarrow \{AC\} = \{A\} \{C\}$
- (4) $I(B,C) \leftrightarrow \{BC\} = \{B\} \{C\}$
- (5) $I(A,BC) \leftrightarrow \{ABC\} = \{A\} \{BC\}$
- (6) $I(B,AC) \leftrightarrow \{ABC\} = \{B\} \{AC\}$
- (7) $I(C,AB) \leftrightarrow \{ABC\} = \{C\} \{AB\}$
- (8) $I(A,B|C) \leftrightarrow \{AB|C\} = \{A|C\} \{B|C\}$
- (9) $I(A,C|B) \leftrightarrow \{AC|B\} = \{A|B\} \{C|B\}$
- (10) $I(B,C|A) \leftrightarrow \{BC|A\} = \{B|A\} \{C|A\}$

Relationships Among Independence Equations

Since each of these independence equations can be either true or not true, there would be 2^{10} or 1,024 possible combinations of

dependence and independence among three random variables if it were not for the fact that some of the independence equations imply others. We can use the relationships among the ten independence equations to eliminate most of the 1,024 possible combinations, leaving twenty-two possible combinations of independence assertions that can exist among three random variables.

The first relationship among the independence equations is that mutual independence implies all of the other types of independence.

$$I(A,B,C) \rightarrow I(A,B), I(A,C), I(B,C), I(A,BC), I(B,AC), I(C,AB), \\ I(A,B|C), I(A,C|B), I(B,C|A)$$

To prove this relationship we start with the equation for mutual independence $\{ABC\} = \{A\} \{B\} \{C\}$. If we integrate both sides of this equation over all possible values of C , we have $\{AB\} = \{A\} \{B\}$. This proves that A and B are independent without regard to C , or $I(A,B)$. Similarly, by integrating over all possible values B and A , we find that $\{AC\} = \{A\} \{C\}$ and $\{BC\} = \{B\} \{C\}$. These two equations are equivalent to $I(A,C)$ and $I(B,C)$. If we substitute these equations into the equation for mutual independence, we have $\{ABC\} = \{A\} \{BC\} = \{B\} \{AC\} = \{C\} \{AB\}$, or $I(A,BC)$, $I(B,AC)$, and $I(C,AB)$. By using these results we show that

$$\{AB|C\} = \frac{\{ABC\}}{\{C\}} = \frac{\{A\}\{B\}\{C\}}{\{C\}} = \{A\}\{B\} = \{A|C\} \{B|C\} \text{ when } \{C\} \neq 0^*.$$

* This is not a restrictive condition since $\{AB|C\}$, $\{A|C\}$, and $\{B|C\}$ are not defined when $\{C\} = 0$.

This is equivalent to $I(A, B|C)$. In exactly the same way we can prove $I(A, C|B)$ and $I(B, C|A)$ when $B \neq 0$ and $A \neq 0$. Thus, mutual independence implies all other types of independence.

The other relationships among the independence equations are listed below. The proofs for these relationships are similar to the one for mutual independence, and they are outlined briefly.

$$1. \quad I(A, B, C) \rightarrow I(A, B), I(A, C), I(B, C), I(A, BC), I(B, AC), I(C, AB), \\ I(A, B|C), I(A, C|B), I(B, C|A)$$

$$2. \quad I(A, BC) \rightarrow I(A, B), I(A, C)$$

$$\text{Proof: } I(A, BC) \rightarrow \{ABC\} = \{A\} \{BC\} \rightarrow \int_c \{ABC\} = \int_c \{A\} \{BC\} \\ \rightarrow \{AB\} = \{A\} \{B\} \rightarrow I(A, B)$$

$$3. \quad I(A, BC) \rightarrow I(A, B|C), I(A, C|B)$$

$$\text{Proof: } I(A, BC) \rightarrow I(A, B), I(A, C) \rightarrow \{AB\} = \{A\} \{B\}, \{AC\} = \{A\} \{C\} \\ \{A\} = \frac{\{AB\}}{\{B\}} = \frac{\{AC\}}{\{C\}} \\ I(A, BC) \rightarrow \{ABC\} = \{A\} \{BC\} = \frac{\{AB\} \{BC\}}{\{B\}} = \frac{\{AC\} \{BC\}}{\{C\}} \rightarrow \\ I(A, B|C), I(A, C|B)$$

$$4. \quad I(A, BC), I(B, C) \rightarrow I(A, B, C)$$

$$\text{Proof: } I(B, C) \rightarrow \{BC\} = \{B\} \{C\} \\ I(A, BC) \rightarrow \{ABC\} = \{A\} \{BC\} = \{A\} \{B\} \{C\} \rightarrow I(A, B, C)$$

$$5. \quad I(A, BC), I(B, AC) \rightarrow I(A, B, C)$$

$$\text{Proof: } I(B, AC) \rightarrow I(B, C) \\ I(A, BC), I(B, C) \rightarrow I(A, B, C)$$

$$6. \quad I(A, BC), I(B, C|A) \rightarrow I(A, B, C)$$

$$\text{Proof: } I(A, BC) \rightarrow I(A, B), I(A, C) \rightarrow \{AB\} = \{A\}\{B\}, \{AC\} = \{A\}\{C\}$$

$$I(B, C|A) \rightarrow \{ABC\} = \frac{\{AB\}\{AC\}}{\{A\}} = \{A\}\{B\}\{C\} \rightarrow I(A, B, C)$$

$$7. \quad I(B, C|A), I(A, B) \rightarrow I(B, AC)$$

$$\text{Proof: } I(A, B) \rightarrow \{AB\} = \{A\}\{B\}$$

$$I(B, C|A) \rightarrow \{ABC\} = \frac{\{AB\}\{AC\}}{\{A\}} = \{B\}\{AC\} \rightarrow I(B, AC)$$

Combinations of Independence Equations

The twenty-two possible combinations of independence equations are shown graphically in Figure 2.7. The remaining 1,002 combinations of independence equations are not possible since they violate one of the relationships above. The twenty-two combinations of independence equations have been labeled with numbers in square brackets at the right of Figure 2.7. The first combination corresponds to mutual independence and the twenty-second combination corresponds to mutual dependence. The other twenty combinations correspond to various intermediate levels of independence among three random variables.

Each of the combinations of independence equations is demonstrated below with a simple example based on flipping four types of coins: fair, biased, thick and magnetized coins. The random variables in each example describe whether a head or tail results from flipping the coins. These random variables are discrete, but it is easy to generalize the examples to include continuous random variables. The twenty-two examples are numbered to match the twenty-two combinations of independence assumptions shown in Figure 2.7.

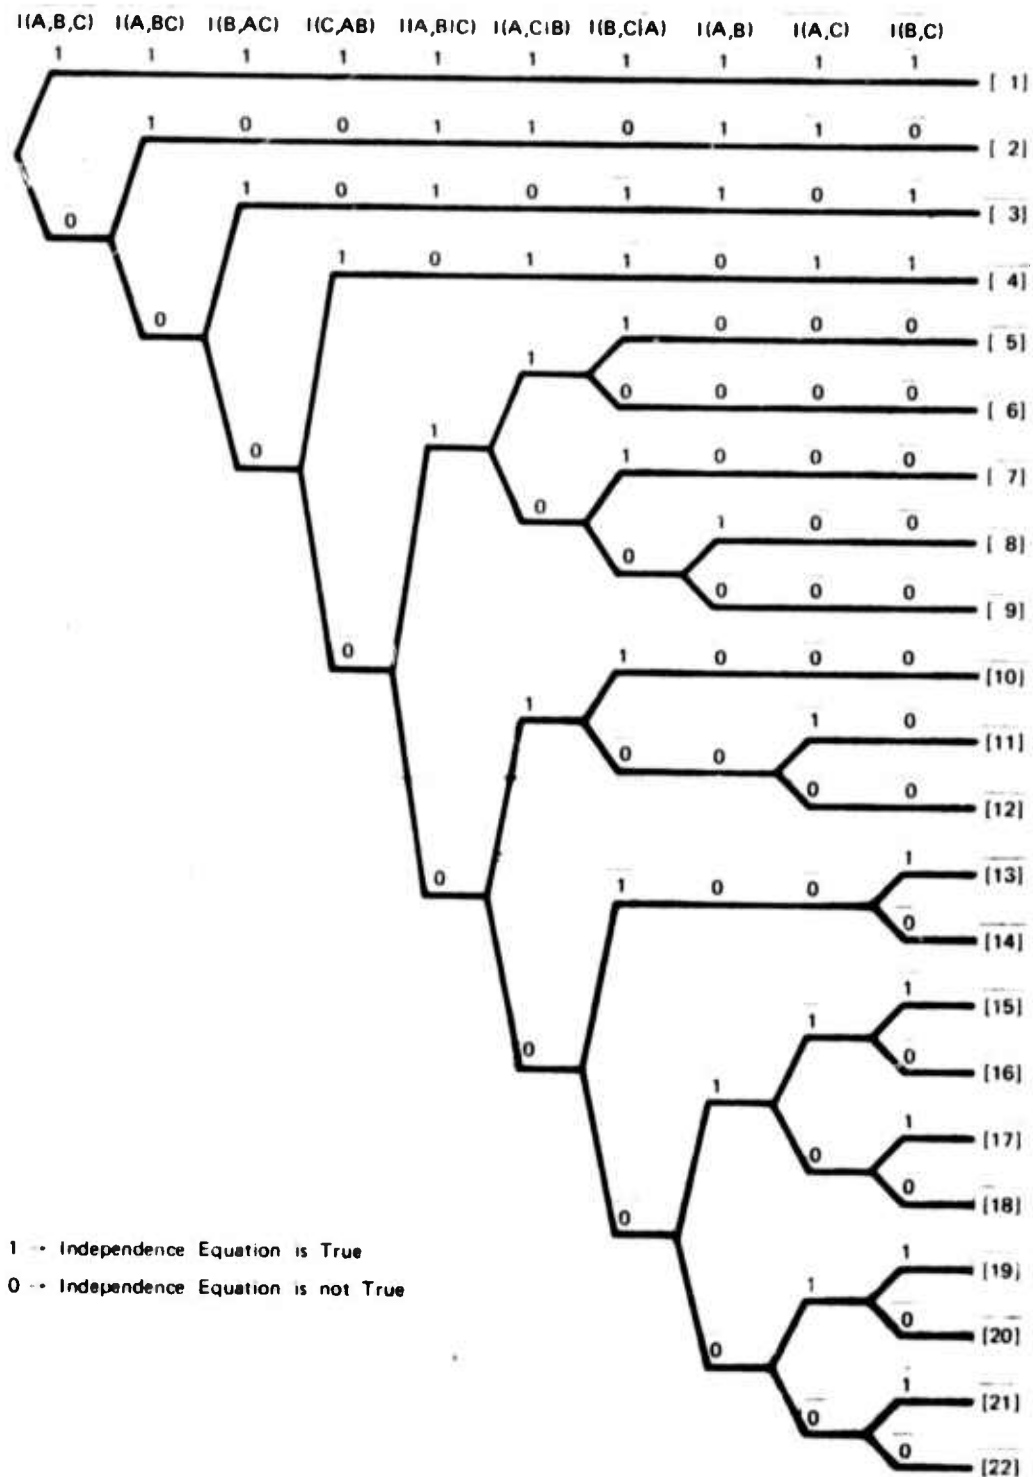


FIGURE 2.7 COMBINATIONS OF INDEPENDENCE EQUATIONS

Example [1]: $I(A,B,C), I(A,BC), I(B,AC), I(C,AB), I(A,B|C),$
 $I(A,C|B), I(B,C|A), I(A,B), I(A,C), I(B,C)$

In this example $A, B,$ and C correspond to the outcomes of flipping three fair coins. The joint probability mass function for $A, B,$ and C is:

		B			
		H	T		
A	H	0.125	0.125	A	H
	T	0.125	0.125		T
		C = H			
		H	T		
		0.125	0.125		
		0.125	0.125		
		C = T			

Example [2]: $\bar{I}(A,B,C), I(A,BC), \bar{I}(B,AC), \bar{I}(C,AB), I(A,B|C),$
 $I(A,C|B), \bar{I}(B,C|A), I(A,B), I(A,C), \bar{I}(B,C)$

In this example A corresponds to the outcome of flipping a fair coin, and B and C correspond to the outcomes of flipping two magnetic coins. The probability that a head will occur when the first magnetic coin is tossed is 50%. However, there is a 60% chance that the second magnetic coin will land with the same side up as the first magnetic coin. Neither of the magnetic coins are affected by the outcome of tossing the fair coin. The conditional probabilities for this example are:

$$\{A=H\} = \{B=H\} = 0.5$$

$$\{C=H|B=H\} = \{C=T|B=T\} = 0.6$$

The joint probability mass function for the three random variables is:

		C		C			
		H	T	H	T		
B	H	0.15	0.10	0.15	0.10	H	
	T	0.10	0.15	0.10	0.15	T	
		A = H		A = T			

Example [3]: $\bar{I}(A,B,C), \bar{I}(A,BC), I(B,AC), \bar{I}(C,AB), I(A,B|C),$
 $\bar{I}(A,C|B), I(B,C|A), I(A,B), \bar{I}(A,C), I(B,C)$

This example is similar to Example [2] with the random variables permuted.

Example [4]: $\bar{I}(A,B,C), \bar{I}(A,BC), \bar{I}(B,AC), I(C,AB), \bar{I}(A,B|C),$
 $I(A,C|B), I(B,C|A), \bar{I}(A,B), I(A,C), I(B,C)$

This example is similar to Example [2] with the random variables permuted.

Example [5]: $\bar{I}(A,B,C), \bar{I}(A,BC), \bar{I}(B,AC), \bar{I}(C,AB), I(A,B|C),$
 $I(A,C|B), I(B,C|A), \bar{I}(A,B), \bar{I}(A,C), \bar{I}(B,C)$

In this example A, B, and C each correspond to the same flip of a fair coin. The conditional probabilities for this example are:

$$\{A=H\} = 0.5$$

$$\{B=H|A=H\} = \{B=T|A=T\} = 1.0$$

$$\{C=H|A=H\} = \{C=T|A=T\} = 1.0$$

The joint probability mass function for the three random variables is:

		B		B	
		H	T	H	T
A	H	0.5	0	0	0
	T	0	0	0	0.5
		C = H		C = T	

Example [6]: $\bar{I}(A,B,C), \bar{I}(A,BC), \bar{I}(B,AC), \bar{I}(C,AB), I(A,B|C),$
 $I(A,C|B), \bar{I}(B,C|A), \bar{I}(A,B), \bar{I}(A,C), \bar{I}(B,C)$

In this example, A corresponds to the outcome of flipping a fair coin. If A is a head, then B and C must also be heads. If A is a tail, then B and C both correspond to the same flip of another fair coin. The conditional probabilities for this example are:

$$\{A=H\} = 0.5$$

$$\{B=H|A=H\} = 1.0, \quad \{B=H|A=T\} = 0.5$$

$$\{C=H|B=H\} = \{C=T|B=T\} = 1.0$$

The joint probability mass function for the three random variables is:

		C				C		
		H	T			H	T	
B	H	0.5	0			0.25	0	H
	T	0	0			0	0.25	T
		A = H				A = T		

Example [7]: $\bar{I}(A,B,C)$, $\bar{I}(A,BC)$, $\bar{I}(B,AC)$, $\bar{I}(C,AB)$, $I(A,B|C)$,
 $\bar{I}(A,C|B)$, $I(B,C|A)$, $\bar{I}(A,B)$, $\bar{I}(A,C)$, $\bar{I}(B,C)$

This example is similar to Example [6] with the random variable permuted.

Example [8]: $\bar{I}(A,B,C)$, $\bar{I}(A,BC)$, $\bar{I}(B,AC)$, $\bar{I}(C,AB)$, $I(A,B|C)$,
 $\bar{I}(A,C|B)$, $\bar{I}(B,C|A)$, $I(A,B)$, $\bar{I}(A,C)$, $\bar{I}(B,C)$

In this example C corresponds to the outcome of flipping a thick coin (a cylinder) that can land on its edge in addition to heads or tails. Possible outcomes for C are heads (H), tails (T), and edge (E). A and B correspond to the outcome of flipping two biased coins, where the bias of each coin depends on the outcome of C. The conditional probabilities for this example are:

$$\begin{aligned} \{A=H|C=H\} &= 0.4, & \{B=H|C=H\} &= 0.6 \\ \{A=H|C=T\} &= 0.5, & \{B=H|C=T\} &= 0.1 \\ \{A=H|C=E\} &= 0.7, & \{B=H|C=E\} &= 0.6 \\ \{C=H\} &= \{C=T\} = 0.4, & \{C=E\} &= 0.2 \end{aligned}$$

The joint probability mass function for the three random variables is:

		B				B				B			
		H	T			H	T			H	T		
A	H	0.096	0.064			0.02	0.18			0.084	0.056	H	
	T	0.144	0.096			0.02	0.18			0.036	0.024	T	A
		C=H				C=T				C=E			

It can be shown that this combination of independence assertions cannot exist for three binary events. Thus it was necessary for C to have three possible outcomes in the example above.

Example [9]: $\bar{I}(A,B,C)$, $\bar{I}(A,BC)$, $\bar{I}(B,AC)$, $\bar{I}(C,AB)$, $I(A,B|C)$,
 $\bar{I}(A,C|B)$, $\bar{I}(B,C|A)$, $\bar{I}(A,B)$, $\bar{I}(A,C)$, $\bar{I}(B,C)$

In this example, C corresponds to the outcome of flipping a fair coin. A and B correspond to the outcomes of flipping two biased coins, where the direction in which the coins are biased depends on C. The conditional probabilities for this example are:

$$\{A=H|C=H\} = \{B=H|C=H\} = 0.6$$

$$\{A=H|C=T\} = \{B=H|C=T\} = 0.4$$

$$\{C=H\} = 0.5$$

The joint probability mass function for the three random variables is:

		B				B			
		H	T			H	T		
A	H	0.18	0.12			0.08	0.12	H	A
	T	0.12	0.08			0.12	0.18		
		C=H				C=T			

Example [10]: $\bar{I}(A,B,C), \bar{I}(A,BC), \bar{I}(B,AC), \bar{I}(C,AB), \bar{I}(A,B|C),$
 $I(A,C|B), I(B,C|A), \bar{I}(A,B), \bar{I}(A,C), \bar{I}(B,C)$

This example is similar to Example [6] with the random variables permuted.

Example [11]: $\bar{I}(A,B,C), \bar{I}(A,BC), \bar{I}(B,AC), \bar{I}(C,AB), \bar{I}(A,B|C),$
 $I(A,C|B), \bar{I}(B,C|A), \bar{I}(A,B), I(A,C), \bar{I}(B,C)$

This example is similar to Example [8] with the random variables permuted.

Example [12]: $\bar{I}(A,B,C), \bar{I}(A,BC), \bar{I}(B,AC), \bar{I}(C,AB), \bar{I}(A,B|C),$
 $I(A,C|B), \bar{I}(B,C|A), \bar{I}(A,B), \bar{I}(A,C), \bar{I}(B,C)$

This example is similar to Example [9] with the random variables permuted.

Example [13]: $\bar{I}(A,B,C), \bar{I}(A,BC), \bar{I}(B,AC), \bar{I}(C,AB), \bar{I}(A,B|C),$
 $\bar{I}(A,C|B), I(B,C|A), \bar{I}(A,B), \bar{I}(A,C), I(B,C)$

This example is similar to Example [8] with the random variables permuted.

Example [14]: $\bar{I}(A,B,C), \bar{I}(A,BC), \bar{I}(B,AC), \bar{I}(C,AB), \bar{I}(A,B|C),$
 $\bar{I}(A,C|B), I(B,C|A), \bar{I}(A,B), \bar{I}(A,C), \bar{I}(B,C)$

This example is similar to Example [9] with the random variables permuted.

Example [15]: $\bar{I}(A,B,C), \bar{I}(A,BC), \bar{I}(B,AC), \bar{I}(C,AB), \bar{I}(A,B|C),$
 $\bar{I}(A,C|B), \bar{I}(B,C|A), I(A,B), I(A,C), I(B,C)$

In this example C corresponds to the outcome of flipping a fair coin. A and B correspond to the outcomes of flipping two magnetized coins, where the direction in which B is magnetized depends on C. The conditional probabilities for this example are:

$$\{C=H\} = \{B=H\} = 0.5$$

$$\{A=H|B=H, C=H\} = \{A=T|B=T, C=H\} = 0.6$$

$$\{A=H|B=H, C=T\} = \{A=T|B=T, C=T\} = 0.4$$

The joint probability mass function for the three random variables is:

		B		B			
		H	T	H	T		
A	H	0.15	0.10	0.10	0.15	H	A
	T	0.10	0.15	0.15	0.10	T	
		C=H		C=T			

Example [16]: $\bar{I}(A,B,C), \bar{I}(A,BC), \bar{I}(B,AC), \bar{I}(C,AB), \bar{I}(A,B|C),$
 $\bar{I}(A,C|B), \bar{I}(B,C|A), I(A,B), I(A,C), \bar{I}(B,C),$

In this example A corresponds to the outcome of flipping a fair coin. B and C correspond to the outcomes of flipping two magnetized coins, where the outcome of A is used to determine how strongly the coins are magnetized. The conditional probabilities for this example are:

$$\{A=H\} = \{B=H\} = 0.5$$

$$\{C=H|B=H, A=H\} = \{C=T|B=T, A=H\} = 0.8$$

$$\{C=H|B=H, A=T\} = \{C=T|B=T, A=T\} = 0.6$$

The joint probability mass function for the three random variables is:

		C				C			
		H	T			H	T		
B	H	0.20	0.05			H	0.15	0.10 <th rowspan="3">H</th>	H
	T	0.05	0.20			T	0.10	0.15	
A=H				A=T					

Example [17]: $\bar{I}(A,B,C), \bar{I}(A,BC), \bar{I}(B,AC), \bar{I}(C,AB), \bar{I}(A,B|C),$
 $\bar{I}(A,C|B), \bar{I}(B,C|A), I(A,B), \bar{I}(A,C), I(B,C)$

This example is similar to Example [16] with the random variables permuted.

Example [18]: $\bar{I}(A,B,C), \bar{I}(A,BC), \bar{I}(B,AC), \bar{I}(C,AB), \bar{I}(A,B|C),$
 $\bar{I}(A,C|B), \bar{I}(B,C|A), I(A,B), \bar{I}(A,C), \bar{I}(B,C)$

In this example A and B correspond to outcomes of flipping two fair coins. C corresponds to the outcome of flipping a biased coin where the amount that the coin is biased depends on A and B. The conditional probabilities for this example are:

$$\{A=H\} = \{B=H\} = 0.5$$

$$\{C=H|A=T, B=T\} = 0.4$$

$$\{C=H|A=H, B=T\} = \{C=H|A=T, B=H\} = 0.6$$

$$\{C=H|A=H, B=H\} = 0.8$$

The joint probability mass function for the three random variables is:

		B				B			
		H	T			H	T		
A	H	0.20	0.15			H	0.05	0.10 <th rowspan="3">H</th>	H
	T	0.15	0.10			T	0.10	0.15	
		C=H				C=T			

Example [19]: $\bar{I}(A,B,C), \bar{I}(A,BC), \bar{I}(B,AC), \bar{I}(C,AB), \bar{I}(A,B|C),$
 $\bar{I}(A,C|B), \bar{I}(B,C|A), \bar{I}(A,B), I(A,C), I(B,C)$

This example is similar to Example [16] with the random variables permuted.

Example [20]: $\bar{I}((A,B,C), \bar{I}(A,BC), \bar{I}(B,AC), \bar{I}(C,AB), \bar{I}(A,B|C),$
 $\bar{I}(A,C|B), \bar{I}(B,C|A), \bar{I}(A,B), I(A,C), \bar{I}(B,C)$

This example is similar to Example [18] with the random variables permuted.

Example [21]: $\bar{I}(A,B,C), \bar{I}(A,BC), \bar{I}(B,AC), \bar{I}(C,AB), \bar{I}(A,B|C),$
 $\bar{I}(A,C|B), \bar{I}(B,C|A), \bar{I}(A,B), \bar{I}(A,C), I(B,C)$

This example is similar to Example [18] with the random variables permuted.

Example [22]: $\bar{I}(A,B,C), \bar{I}(A,BC), \bar{I}(B,AC), \bar{I}(C,AB), \bar{I}(A,B|C),$
 $\bar{I}(A,C|B), \bar{I}(B,C|A), \bar{I}(A,B), \bar{I}(A,C), \bar{I}(B,C)$

In this example, A , B , and C , correspond to the outcome of flipping three magnetic coins. There is a 50% chance that the first coin flipped will come up heads, there is a 60% chance that the second coin flipped lands with the same side up as the first. If the first two coins are both heads or both tails, there is an 80% chance that the third coin flipped will have the same outcome; otherwise the probability of heads on the third flip is 50%. The conditional probabilities for this example are:

$$\{A=H\} = 0.5$$

$$\{B=H|A=H\} = \{B=T|A=T\} = 0.6$$

$$\{C=H|A=H,B=H\} = \{C=T|A=T,B=T\} = 0.8$$

$$\{C=H|A=H,B=T\} = \{C=H|A=T,B=H\} = 0.5$$

The joint probability mass function for the three random variables is:

		B				B			
		H	T			H	T		
A	H	0.24	0.10			0.06	0.10	H	A
	T	0.10	0.06			0.10	0.24		
		C=H				C=T			

An Additional Relationship Among the Independence Equations

Some of the combinations of independence equations require the joint distribution for the three random variables to contain several zeros. This means that certain combinations of outcomes for the three random variables are not possible even though each of the random variables can individually assume the same outcomes. In Example [5], A, B, and C can each be heads or tails, but it is not possible for one to be a head when another is a tail.

If we assume that the joint distribution cannot equal zero when the individual marginal distributions are not zero, it is possible to prove an additional relationship among the ten independence equations.

$$I(A,B|C), I(A,C|B) \rightarrow I(A,BC) \text{ if } \{BC\} \neq 0$$

The assumption that $\{BC\}$ is not zero allows us to divide by this quantity in the following proof:

$$I(A,B|C), I(A,C|B) \rightarrow \{ABC\} = \frac{\{AC\}\{BC\}}{\{C\}} = \frac{\{AB\}\{BC\}}{\{B\}}$$

$$\rightarrow \{AC\}\{B\} = \{AB\}\{C\} \rightarrow \int_c \{AC\}\{B\} = \int_c \{AB\}\{C\}$$

$$\rightarrow \{A\}\{B\} = \{AB\} \rightarrow I(A,B)$$

$$I(A,C|B), I(A,B) \rightarrow I(A,BC)$$

If this relationship is added to the seven discussed above, it can be used to eliminate four of the twenty-two possible combinations of independence equations. The four combinations that are eliminated are numbered [5], [6], [7], and [10] in Figure 2.7. The remaining eighteen combinations of independence equations are still possible.

Encoding the Twenty-Two Combinations of Independence Equations

A subjects' state of information about several uncertain quantities can be represented by a wide variety of possible combinations of independence equations, even when the problem contains as few as three random variables. One of the principal motivations for assuming that random variables are independent, is to limit the amount of probability encoding required to specify the joint probability distribution for all of the random variables. Although there are many possible combinations of independence equations, the degree of difficulty associated with assessing the uncertainties necessary to specify the joint distribution can be determined by some very simple properties of the independence equations that are assumed to be true. We can place the ten possible independence equations for three

random variables, in the four categories as shown below:

Category 1	$I(A,B,C) \leftrightarrow \{ABC\} = \{A\} \{B\} \{C\}$
	$I(A,BC) \leftrightarrow \{ABC\} = \{A\} \{BC\}$
Category 2	$I(B,AC) \leftrightarrow \{ABC\} = \{B\} \{AC\}$
	$I(C,AB) \leftrightarrow \{ABC\} = \{C\} \{AB\}$
	$I(A,B C) \leftrightarrow \{AB C\} = \{A C\} \{B C\}$
Category 3	$I(A,C B) \leftrightarrow \{AC B\} = \{A B\} \{C B\}$
	$I(B,C A) \leftrightarrow \{BC A\} = \{B A\} \{C A\}$
	$I(A,B) \leftrightarrow \{AB\} = \{A\} \{B\}$
Category 4	$I(A,C) \leftrightarrow \{AC\} = \{A\} \{C\}$
	$I(B,C) \leftrightarrow \{BC\} = \{B\} \{C\}$

The degree of difficulty associated with assessing the probabilities needed to specify the joint distribution for A, B, and C depends only on which categories of independence assumptions contain equations that are assumed to be true. As a measure of the degree of difficulty associated with the assessment problem, we will assume that A, B, and C are each discrete random variables with probability mass functions that contain n possible outcomes; we will then determine the minimum number of probabilities required to specify the joint mass function.

If the three random variables are assumed to be mutually independent (Category 1), then the joint probability density function can be determined by assessing the three marginal distributions and multiplying them together. To specify each of the probability mass functions, we would need to assess n probabilities. Therefore, to determine the joint probability distribution for A, B, and C, we would need to assess $3n$ probabilities.

If our state of information about the three random variables is such that we can not assume mutual independence, but can assume that one of the independence equations in Category 2 is true, we can determine the joint probability density function by assessing one of the ~~marginal~~ distributions and the joint distribution for the two remaining random variables. We need to assess n probabilities to determine the marginal distribution, and n^2 probabilities to determine the joint distribution for two random variables. Thus, when we can assume that one of the independence equations in Category 2 is true, but that the independence equation in Category 1 is not true, we would need to assess $(n^2 + n)$ probabilities. This situation occurs for the combinations of independence equations numbered [2], [3], and [4] in Figure 2.7.

If our state of information about the three random variables is such that we cannot assume that the independence equations in Categories 1 and 2 are true, but that one of the independence equations in Category 3 is true, we can determine the joint probability density function by assessing the marginal distribution from one of the three random variables and the conditional distributions for the other two random variables, given the first. In this situation, we would need to assess n probabilities for the marginal distribution, and n^2 probabilities for each of the two conditional distributions. Thus, when the independence equations in Categories 1 and 2 are not true, but one of the independence equations in Category 3 is true, we will need to assess $(2n^2 + n)$ probabilities. This situation occurs for the combinations of independence equations numbered [5] through [14] in Figure 2.7.

If our state of information about the three random variables is such that none of the independence equations in Categories 1, 2, and 3 are true, but one of the independence equations in Category 4 is true, then we can determine the joint density function by assessing two marginal distributions and the conditional distribution for the third random variable, given the first two. However, to assess the conditional distribution, we would need to assess n^3 probabilities.

Since we can assess the joint distribution for all three random variables with n^3 probabilities, we can minimize the number of probabilities assessed by doing so. We could also assess the joint distribution for all three random variables by assessing the n^3 probabilities in the case where none of the independent equations are true. Thus, when none of the independence equations in Categories 1, 2, and 3 are true, we will need to assess at least n^3 probabilities.

The number of probabilities that must be assessed to determine the joint distribution for all three random variables is summarized in the following table as a function of the categories of independence equations.

Categories of Independence Equations	Number of Probabilities Assessed
1	$3n$
2	$n^2 + n$
3	$2n^2 + n$
4	n^3

The number of probabilities that must be assessed to determine the joint distribution is a rough measure of the degree of difficulties associated with the encoding process. For a three variable problem with continuous probability density functions or discrete mass functions with many possible outcomes, cases corresponding to large values of n , , mutual

independence (Category 1) or partial independence (Categories 2 and 3) can be powerful simplifying assumptions. Where more than three variables are involved the independence assumptions are even more powerful.

2.2.2 Coalescence

Coalescence is a process of eliminating redundant portions of decision trees. Coalescence is the key to solving large complex decision problems that would require a prohibitive amount of probability encoding and processing if the redundancies were not eliminated. The redundancies occur, and coalescence becomes possible, when two or more sections of a decision tree contain exactly the same structure and numerical data (probabilities and values). However, for two sections of a decision tree to have the same probabilities, the probabilities must correspond to random variables which are independent of other random variables and decisions represented in the decision tree. Thus, the use of coalescence to simplify a decision tree requires an understanding to the independencies that exist among the random variables.

Although a decision tree specifies the probabilistic structure of a decision problem, it does not do a good job of describing the independencies among random variables and decisions. If two random variables in a decision tree

are independent the probabilities associated with one random variable are different in portions of the tree that correspond to different outcomes for the other random variable.. The necessity to search the tree to discover dependencies makes it very difficult to visualize or alter the independence assumptions in the decision tree formed.

Influence Diagrams

To overcome this difficulty, a procedure has been developed for describing the dependencies among random variables and decisions in the form of an "influence diagram." An influence diagram can be used to specify and visualize the probabilistic dependencies in a decision analysis, and to determine the degree of coalescence that is possible in the corresponding decision tree. Using influence diagrams an analyst can take full advantage of coalescence to minimize both the assessment of probabilities and the amount of computation required to solve the decision tree.

Figure 2.8 shows how influence diagrams represent the dependencies among random variables and decisions. A random variable is represented by a circle containing its name or number, and a decision is represented by a square containing its name or number. An arrow pointing from random variable A to random variable B means that the outcome of A can influence the probabilities associated with B. An arrow



THE PROBABILITIES ASSOCIATED WITH RANDOM VARIABLE B DEPENDS ON THE OUTCOME OF RANDOM VARIABLE A



THE PROBABILITY OF RANDOM VARIABLE D DEPENDS ON DECISION C.



THE DECISION MAKER KNOWS THE OUTCOME OF RANDOM VARIABLE E WHEN DECISION F IS MADE

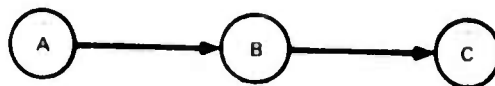


THE DECISION MAKER KNOWS DECISION G WHEN DECISION H IS MADE.

FIGURE 2.8 DEFINITIONS USED IN INFLUENCE DIAGRAMS

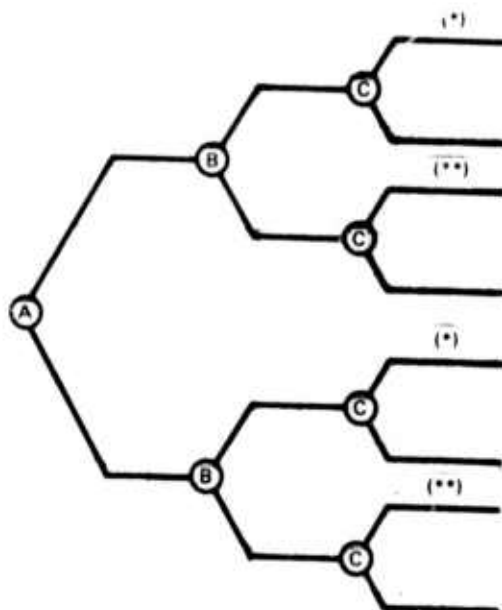
pointing to a decision from either another decision or a random variable means that the decision is made with the knowledge of the outcome of the other decision or random variable. A connected set of squares and circles is called an influence diagram because it shows how random variables and decisions influence each other.

Figure 2.9 shows a simple example of an influence diagram in which all the nodes represent random variables. In this example A influences B, and B influences C. If we do not know the outcome of B, the probabilities associated with C depends on the outcome of A. However, if we know the outcome of B, the probabilities associated with C do not depend on the outcome of A. In other words, A affects the probabilities associated with C through its influence on B but once we know the outcome of B the outcome of A is irrelevant. Figure 2.9 also shows how an influence diagram is translated into a decision tree. In the tree at the bottom of Figure 2.9 the probabilities associated with C do not depend on the outcome of A. Thus, the same probability is attached to the corresponding branches of node C that follow the same outcome for B.



C DEPENDS DIRECTLY ON B AND INDIRECTLY ON A. THE PROBABILITIES ASSOCIATED WITH C DEPEND ON THE OUTCOMES OF A AND B. HOWEVER, GIVEN THE OUTCOME OF B, THE PROBABILITIES ASSOCIATED WITH C DO NOT DEPEND ON THE OUTCOME OF A.

THE CORRESPONDING DECISION TREE IS:



THE PROBABILITIES ASSOCIATED WITH C AT THE TWO BRANCHES MARKED (*) ARE THE SAME. SIMILARLY, THE PROBABILITIES ASSOCIATED WITH C AT THE TWO BRANCHES MARKED (**) ARE THE SAME.

FIGURE 2.9 AN EXAMPLE OF AN INFLUENCE DIAGRAM

Since some of the branches in this probability tree contain the same information, they can be combined, or "coalesced". The coalesced decision tree for this example is shown in Figure 2.10. It contains the same information as the decision tree in Figure 2.9 but has fewer nodes and branches. If the process of coalescence is carried out at the time the decision tree is generated, the size of the model can be greatly reduced. Computer programs have been developed that translate directly from influence diagrams to coalesced decision trees.

In a typical application of influence diagrams, a decision tree with millions of nodes was reduced to one with fewer than a thousand nodes. The influence diagram for this application is shown in Figure 2.11. The corresponding decision tree contains a complex set of interconnected branches. The left hand portion of the decision tree is nearly conventional with the number of branches growing geometrically for each additional state or decision variable; however, the righthand portion coalesces rapidly to relatively few nodes near the end of the tree.

Number of Required Probability Assessments

In addition to providing a way to visualize the dependencies among many uncertain events, and drastically reducing the size of the decision tree, influence diagrams also show the number of probability assessments that are required by

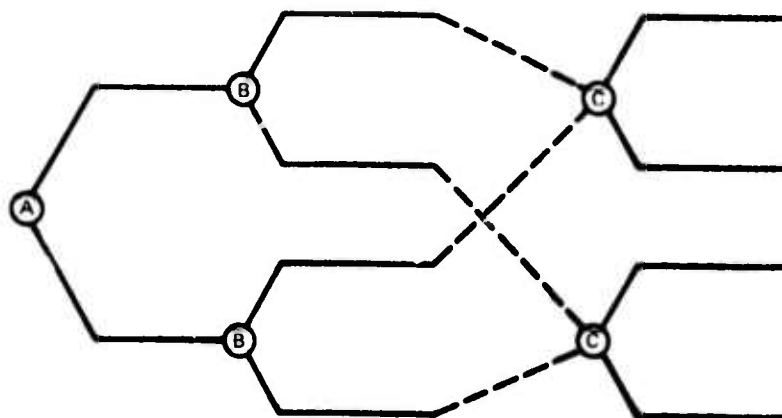


FIGURE 2.10 AN EXAMPLE OF A COALESCED DECISION TREE

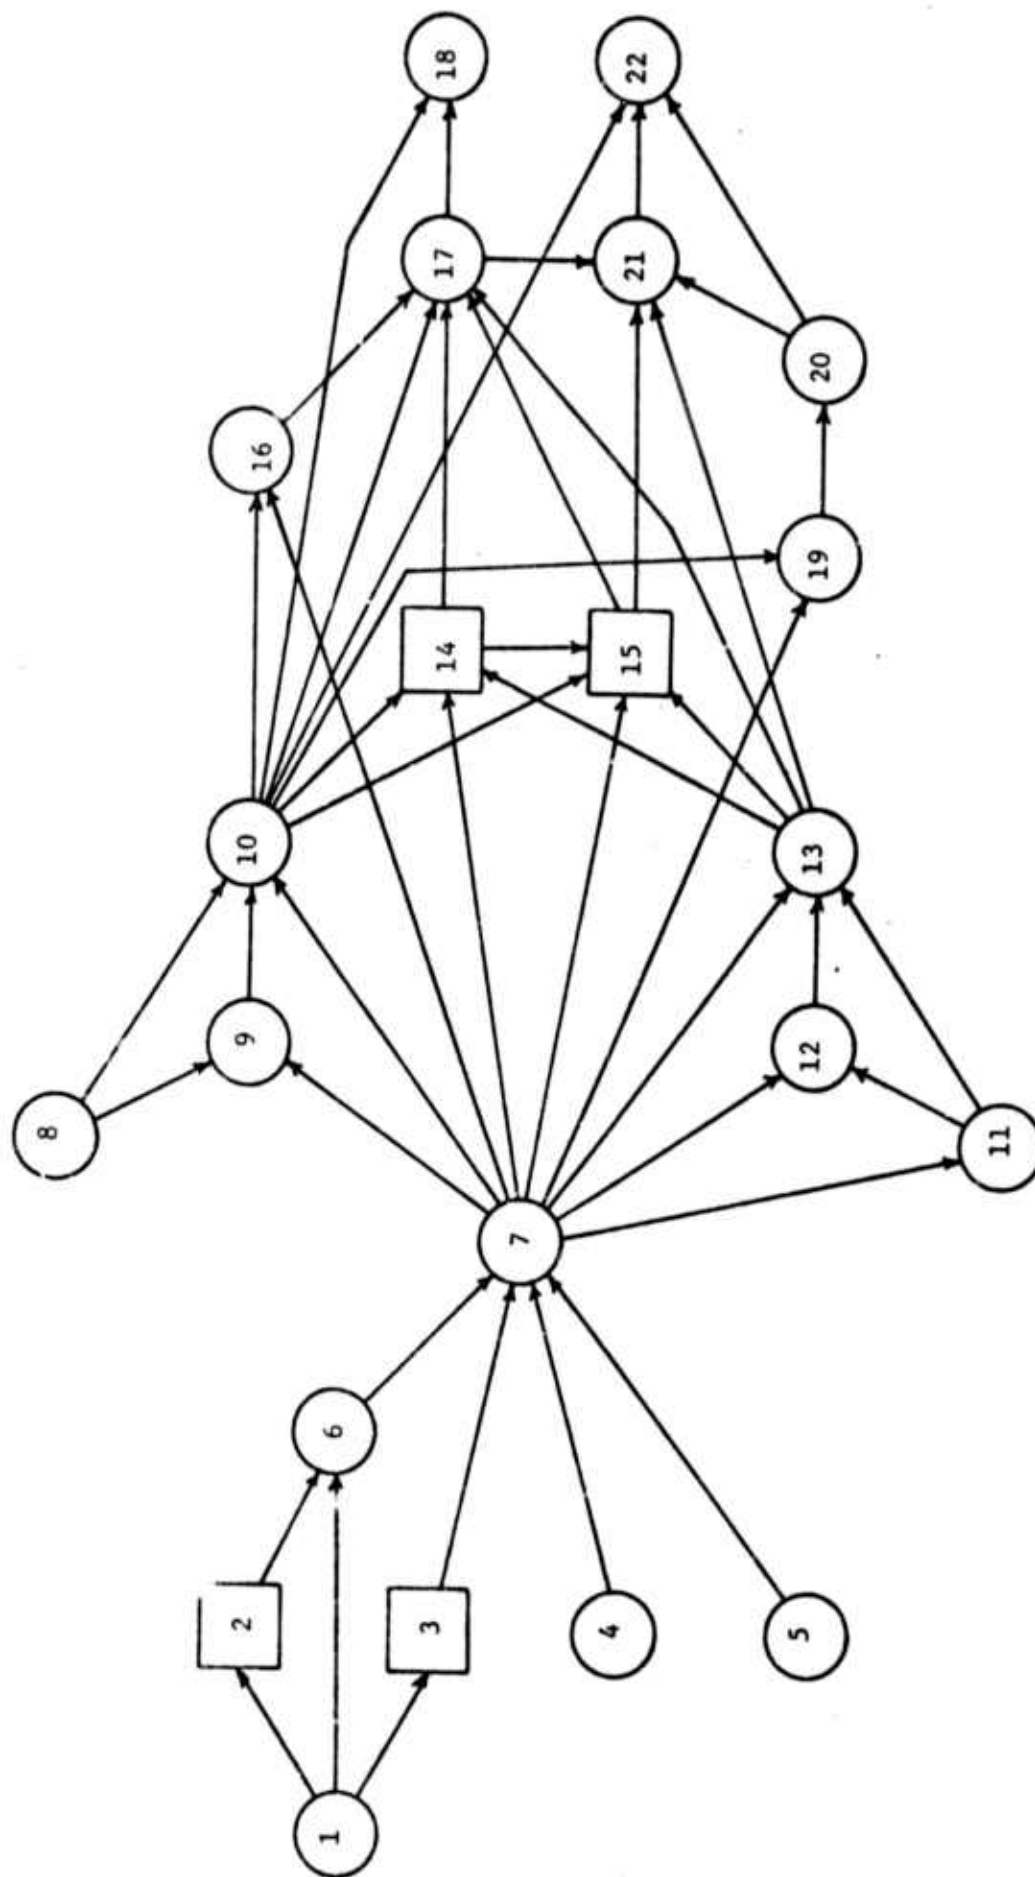


Figure 2.11: AN APPLICATION OF AN INFLUENCE DIAGRAM

the model. If each random variable has N different outcomes, and a particular random variable is directly influenced by M other events or decisions, then it will be necessary to assess M^N probability distributions in order to describe the uncertainty associated with the random variable.

Direction of Arrows in Influence Diagrams

The direction of the arrows in an influence diagram is significant, and in general they cannot be reversed without it changing the independence assumptions implied by the diagram, even though dependence is inherently non-directional. The arrows in the influence diagram cannot be reversed because they indicate both the independencies among random variables and the state of information for which the independence assumptions are true.

For example, the influence diagram in Figure 2.9 means that C is independent of A given B . However, if we modify the influence diagram in Figure 2.9 by reversing one of the arrows, as shown in Figure 2.12 this statement is no longer true. The influence diagram in Figure 2.9 allows us to write

$$\{AC|B\} = \{A|B\} \{C|B\}$$

On the other hand, the influence diagram in Figure 2.12 allows us to write

$$\{AC\} = \{A\} \{C\}$$

Figure 2.12

ANOTHER EXAMPLE OF AN INFLUENCE DIAGRAM



The arrows in the influence diagram in Figure 2.9 show that we must include the knowledge of the random variable B as part of our state of information when asserting that A and C are independent. However, the arrows in the influence diagram in Figure 2.12 show that the independence between A and C occurs before the outcome of random variable B is included in our state of information.

Generating a Decision Tree from an Influence Diagram

To generalize these observations we need to formalize the procedure for translating from an influence diagram to a decision tree. We require that influence diagrams contain no loops, so there must be at least one node in the influence diagram that does not have any arrows pointing toward it. If there is only one such node, then the random variable or decision corresponding to this node is placed at the beginning of the decision tree. If there are several nodes which do not have arrows pointing to them and some of them represent decisions, we can place any one of these decisions at the beginning of the decision tree. If none of the nodes that have no arrows pointing toward them represent decisions, we can place any one of the corresponding random variables at the beginning of the decision tree. Once we have determined the first node in the decision tree, we can remove the corresponding node from the influence diagram along

with all the arrows that leave this node.

After we determine the first node in the decision tree and remove the corresponding node from the influence diagram, we can use the same procedure to determine the second level of the decision tree; the second level consists of all the nodes attached to branches leaving the first node of the tree. The reduced influence diagram (the influence diagram that remains after the node corresponding to the first random variable or decision in the decision tree is removed) must contain at least one node which does not have any arrows pointing toward it. If there is only one such node, the decision or random variable corresponding to this node is placed in the second level of the decision tree. If the reduced influence diagram contains more than one node with no arrows pointing toward it, and some of those nodes correspond to decisions, any one of those decisions can be placed at the second level of the decision tree. However, if all of the nodes that have no arrows pointing toward them represent random variables, then any one of these random variables can be placed at the second level of the decision tree. Once a decision or random variable has been placed at the second level of the decision tree, the corresponding node is removed

from the influence diagram, The same procedure is then used to determine the decision or random variable to be placed at the third level of the decision tree.

For example, if this procedure were applied to the influence diagram shown in Figure 2.13 one of the many possible orders in which the nodes could be placed in a decision tree is shown by the numbers in Figure 2.13. Alternatively, the same influence diagram can be translated into a decision tree with the nodes 2 and 3 reversed, or with nodes 6 and 7 reversed. In fact, there are many other possible ways that the influence diagram in Figure 2.13 can be translated into a decision tree. However, all of the resulting decision trees, though different, would produce the same sets of optimum decisions and same expected values at the beginning of the tree.

Further research is needed before we will have a complete understanding of the meaning of influence diagrams. It is possible to draw influence diagrams that cannot be translated directly into a decision tree using the rules given above. For example, consider the influence diagram shown in Figure 2.14a. If we were to follow the rules given in the preceding paragraphs, we would find that the decision tree corresponding to this influence diagram is the one shown

Figure 2.13

DETERMINING THE ORDER IN WHICH RANDOM VARIABLES
AND DECISIONS ARE PLACED IN A DECISION TREE

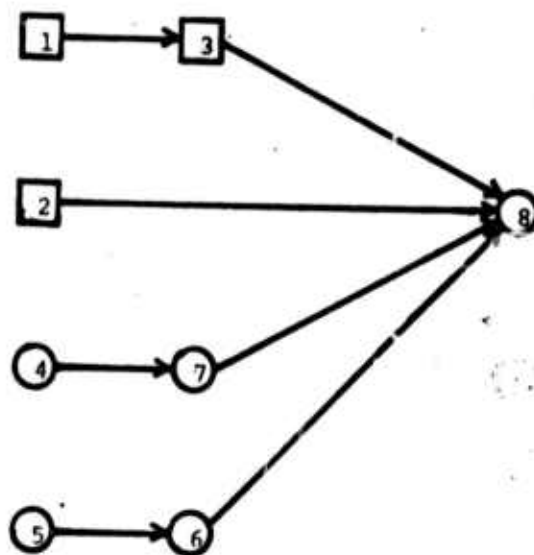


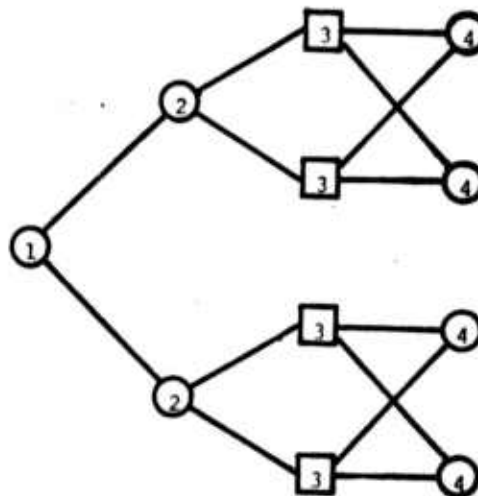
Figure 2.14

AN INFLUENCE DIAGRAM THAT CANNOT BE TRANSLATED
DIRECTLY INTO A DECISION TREE

(a) An Influence Diagram



(b) An Incorrect Coalesced Decision Tree



in Figure 2.14b. In this case, the random variables and decisions must be placed in a decision tree in the order shown. The difficulty with the decision tree shown in Figure 2.14b is that the decisions taken at node 3 may depend on the outcome of the random variable corresponding to node 1, even though the influence diagram in Figure 2.14a shows that node 3 does not depend directly on node 1. On the other hand, if we attempt to coalesce the decision nodes in Figure 2.14b so that they do not depend on the outcome of the random variable corresponding to node 1, then we will not know the state of this random variable when we reach the nodes at the end of the decision tree.

The proper way to generate a decision tree from the influence diagram shown in Figure 2.14a is to use Bayes' rule to reverse the order of random variables 1 and 2 and then place random variable 1 after decision 3 in the tree. In the next phase of our research, we hope to develop general algorithms for translating any influence diagram into a consistent decision tree.

2.2.3 The Value of Decision-Dependent Clairvoyance

The idea of valuing perfect information has appeared in many treatments of decision making under uncertainty. Most often the example being treated represents a simple hypothetical situation. The informational structure that is being captured in probability assignments is straightforward and the assumptions regarding the probabilistic structure, such as types of independence among both controlled and uncontrolled variables, are implicit in the problem statement. However, the correct computation of the value of information can be elusive on both conceptual and numerical bases. The concept of clairvoyance will lead us to the construction of detailed information models and to the exploration of their precise interpretation and use.

The Primary Decision

For illustration, let us play the role of analysts for a space mission designed to land a remotely controlled experimental apparatus on the surface of Mars. We have thoroughly analyzed the mission and have summarized our total state of information by assigning a 0.6 probability that the Mars mission will be successful and a corresponding 0.4 probability that it will be a failure. Also we have analyzed the values to be derived from the mission and have put them in monetary units, millions-of-dollars, for example. Let's assume that the value of a successful Mars mission is 50 units

and the value of an unsuccessful one is 10 units. A positive value might be attributed to a failure because attempting the mission has important social value and even a failure will provide knowledge for a better design on the next attempt.

Unexpectedly, several months before launch another nation announces that it will attempt a similar mission to Venus in about one year. Because of the competitive nature of the space race and the important foreign policy implications of technological leadership, we realize that the value of changing our destination and successfully landing on Venus would be quite high. On the other hand, if we attempt to land on Venus and fail we would look foolish for diverting the program, and in any case we would set back the timetable on our extensive Martian exploration program at least two years. When all of these factors are evaluated we find that a successful landing on Venus is worth 100 units and a failure costs 10 units. To our surprise, when we check the feasibility of diverting the mission we find that because of modular design only a few important but thoroughly tested components of the landing system need to be changed, and the mission engineers assign a 0.6 probability of success regardless of destination.

From these assessments, we can lay out the primary decision tree of Figure 2.15. Along each outcome branch emanating from a chance node we have written the conditional probability of following that path, and near each node we have written the value, either assigned or derived, of being at the point in the program represented by that node.

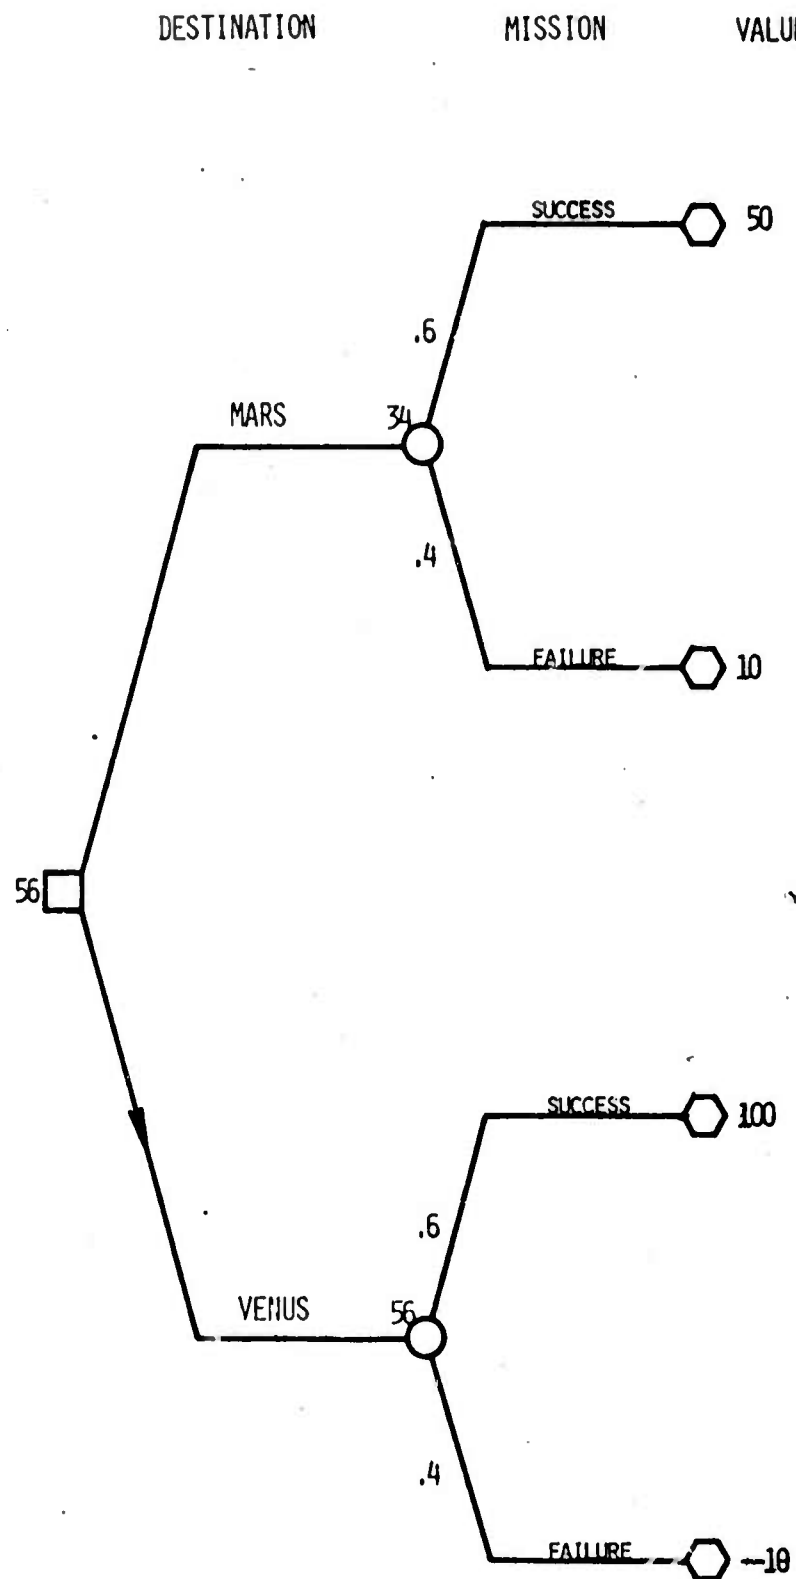


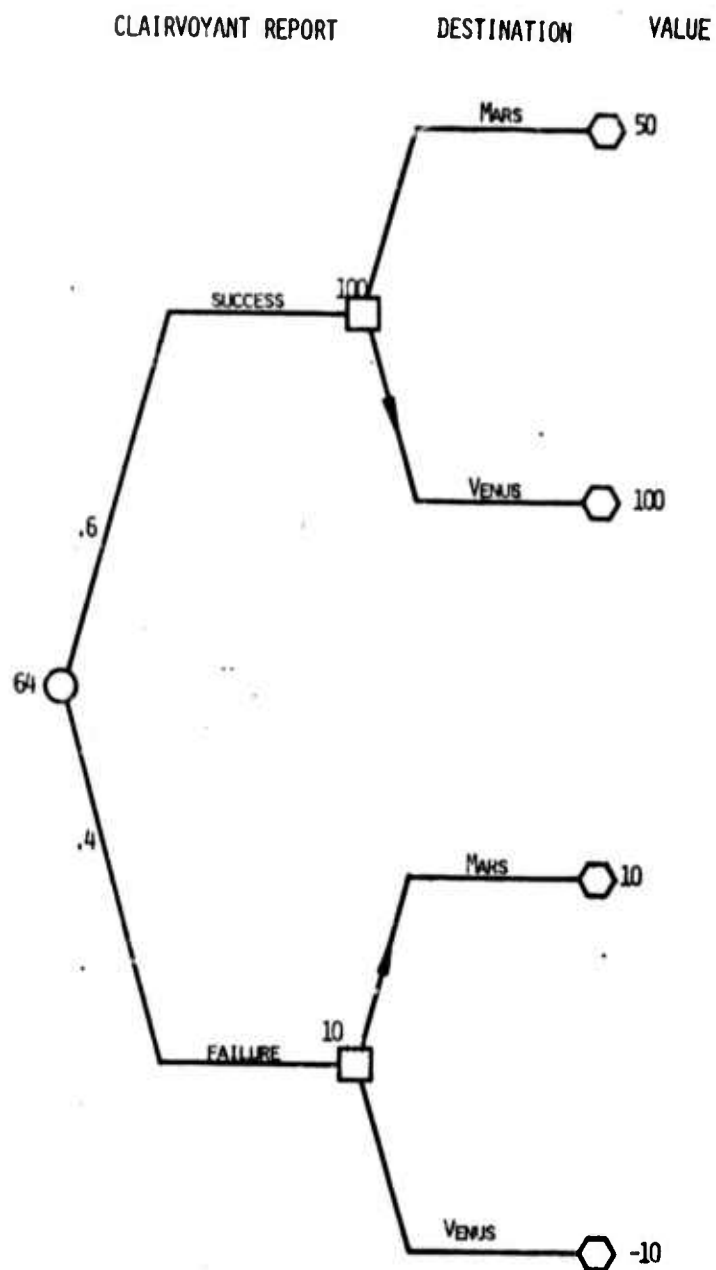
FIGURE 2.15: THE PRIMARY DECISION TREE

We see that the expected value of going to Mars is 34, while the expected value of going to Venus is 56. Thus, in order to maximize the expected value of the mission we decide to go to Venus.

Value of Perfect Information

We might wish to use the decision tree to investigate the possibility of gathering new information before we make the final decision. To do this we can use the value of perfect information as an upper bound for the value of less complete information gathering program. Most analysts, when presented with Figure 2.15 and asked to derive the value of perfect information, reverse the order of decision and chance nodes in Figure 2.15 to produce the tree shown in Figure 2.16. With the latter tree we learn first whether the mission will succeed or fail and then we decide on the destination. If we know the mission will succeed, we send it to Venus for 100 units of value, and, if it will fail, we send it to Mars for 10 units of value. Using the original probability of success (0.6) as the probability that the information will predict a success, we obtain an expected value of 64 units with perfect information. Subtracting 56 units for the value of the primary decision problem, we obtain 8 units for the value of perfect information.

What might be wrong with this approach? Suppose that perfect information revealed that the mission would succeed on Mars but fail on Venus, or vice-versa. These possibilities do not appear in Figure 2.16. To correct this omission we might draw a new tree for the value



with perfect information as illustrated in Figure 2.17. First we learn one of four possible predictions consisting of the four combinations of success or failure on Mars and Venus. Then we make the best decisions given this information, as indicated in the decision tree.

In order to assign probabilities one might reason that, since landings on Mars or Venus appear in separate portions of the primary tree of Figure 2.15, the events must be independent and the probabilities should be multiplied as shown in Figure 2.17. This will yield a value of 73.6 with perfect information, and subtracting the 56 unit value of the primary decision it yields a 17.6 value of perfect information.

However, the independence assumption must be questioned. Since we can only send the mission to a single destination, might the events be mutually exclusive? If we simply try to assign the probabilities directly we are tempted to phrase confusing questions like "What is the probability we will succeed on both Mars and Venus?" or "If we learned we had failed on Mars what probability would we assign to success on Venus?" The trouble stems from the fact that we have only one rocket and it is difficult to consider sending it to both destinations simultaneously, but this consideration seems to be necessary in order to assign the required probabilities. Perhaps

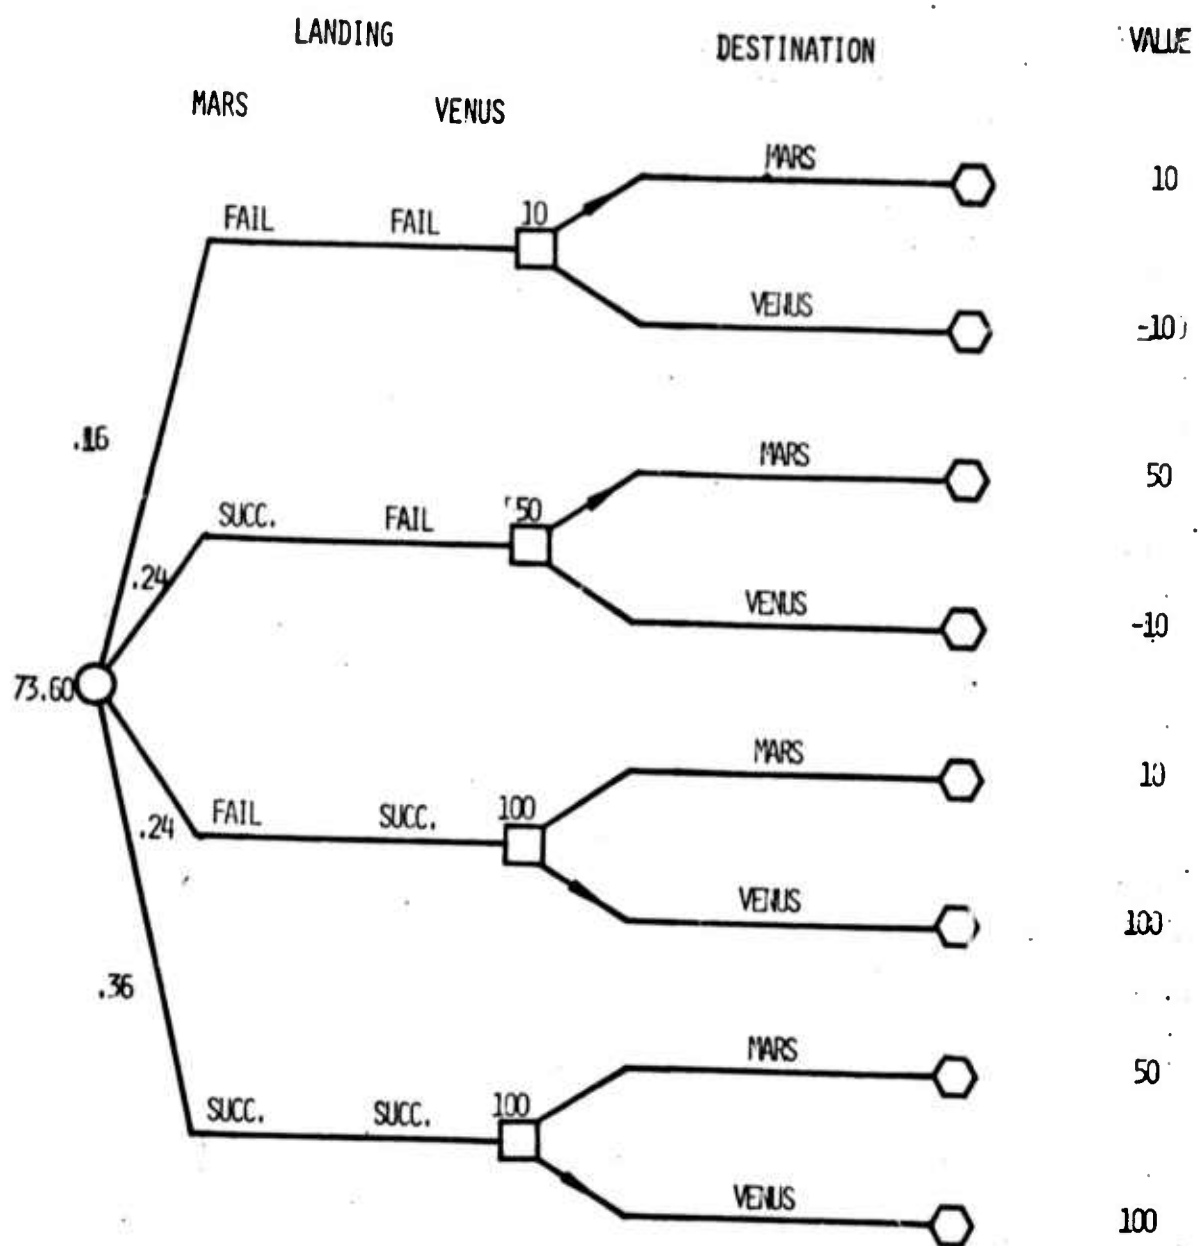


FIGURE 2.17: MORE COMPLEX TREE FOR DETERMINING THE VALUE WITH PERFECT INFORMATION

we could retreat to a "classical" interpretation in which we construct many "identical" hypothetical worlds where some rockets are sent to Mars and others are sent to Venus. With much thought and careful phrasing we might arrive at clear questions with useful interpretations. But we might still wonder if we had come up with a valid assessment.

The Concept of Clairvoyance

These confusing questions of probability assessment can be resolved with the introduction of the clairvoyant, a hypothetical character who knows all and who can answer any well-specified question about any uncertainty. Of course, we shall never really be able to obtain answers from him, but our probability assignments for his possible answers will provide the key to probabilistic structuring.

In the example at hand, if we were to ask the clairvoyant whether the mission will succeed or fail, he might respond by saying that the result could depend on where you sent it. Thus, we would be led to asking him two such questions, one for each destination. To make our questions precise we might draw up the questionnaire of Figure 2.18. Presuming that the clairvoyant is satisfied with our definitions of success and failure, he could answer by checking one box in each row corresponding to the outcome he foretells for each destination choice.

Before we engage the clairvoyant, we wish to calculate the value of his services in monetary units, the value of clairvoyance. Since

DESTINATION	MISSION	
	FAILURE	SUCCESS
MARS	<input type="checkbox"/>	<input type="checkbox"/>
VENUS	<input type="checkbox"/>	<input type="checkbox"/>

CHECK ONE BOX IN EACH ROW

FIGURE 2.18: CLAIRVOYANT'S REPORT FORM

the clairvoyant has two possible answers for each of the two questions there are four possible reports for Mars and Venus: failure, failure; failure, success; success, failure; success, success. We now must assign probabilities to these reports. A possible probability assignment, compatible with our original assignments of Figure 2.15, is illustrated in Figure 2.19. This distribution implies dependence between our knowledge of the clairvoyant's two answers. For example, if he were to answer success on Mars, we would then assign a $0.554/0.6 = 0.923$ probability that he would also answer success on Venus.

Philosophically, the important aspect of this formulation is that we are assigning probabilities to events that could occur immediately, when the clairvoyant reveals his answers. Also, we have avoided the awkward considerations of sending our single spacecraft simultaneously to both planets or of generating hypothetical universes.

We now apply this probability assignment by constructing the decision tree of Figure 2.20. The initial chance node represents the clairvoyant's revelation of one of the four possible reports, each indicated by the abbreviated report form on one of the following branches. The probabilities of Figure 2.20 are assigned to these reports. Following each report we must make the best decision using the values from Figure 2.15. Having made the decisions indicated

		VENUS		
		FAILURE	SUCCESS	
MARS	FAILURE	.354	.046	.4
	SUCCESS	.046	.554	.6
		.4	.6	

FIGURE 2.19: JOINT DISTRIBUTION FOR CLAIRVOYANT'S ANSWERS

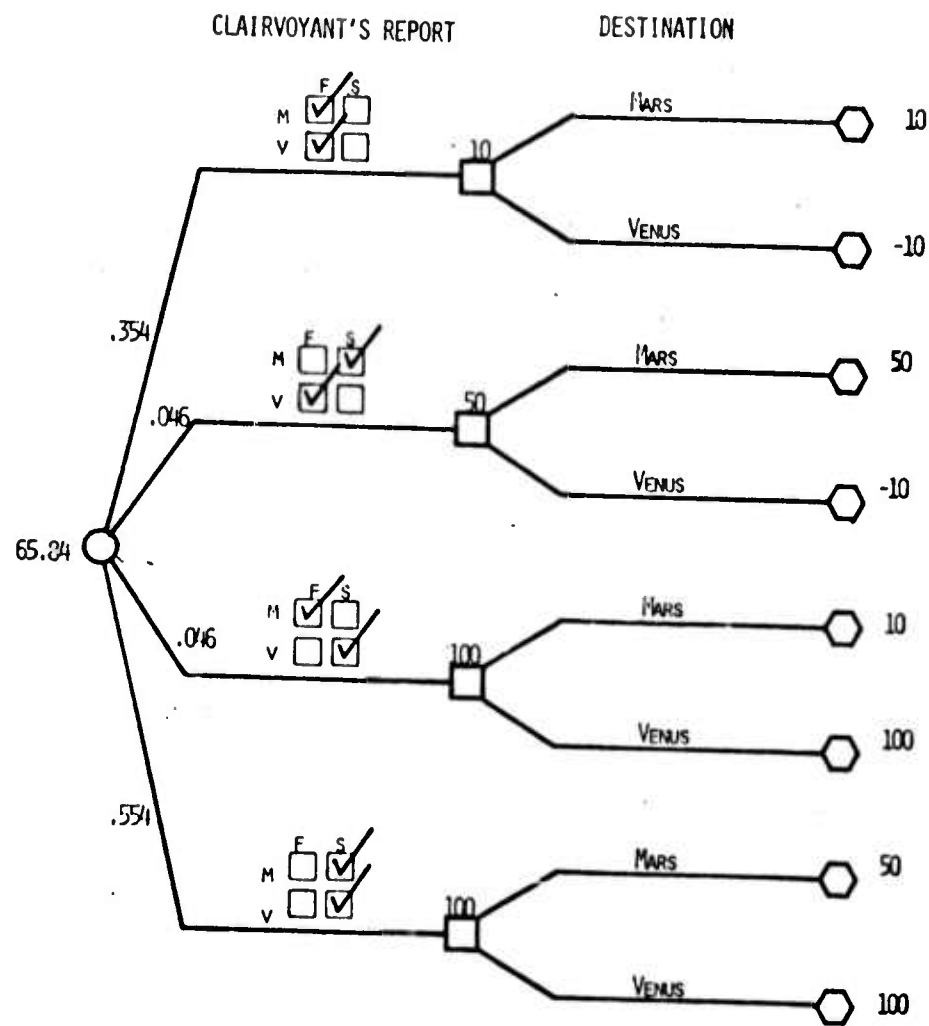


FIGURE 2.20: DECISION TREE FOR DETERMINING THE VALUE OF COMPLETE CLAIRVOYANCE

by the arrows, we find that the expected value with clairvoyance-- but before the clairvoyant reveals his answer--is 65.84. Subtracting the 56 unit value of the primary decision (without clairvoyance) yield a value of clairvoyance of 9.84.

Practical Probability Assignment

It would be rare for experts to think in terms of joint probability distributions such as those of Figure 2.19. Normally, the experts will have technical information organized in a way that is meaningful to them, and it is desirable to construct the probability model in their terms. A simple version of such a model is illustrated by the probability assignment model of Figure 2.21. Here the first question asked of the clairvoyant is, "Will the launch system, which is common to both destinations, work?" The expert has assigned 0.65 to the answer "WORK". The next two questions depend on the destination. The first is "If we launch successfully and send the spacecraft to Mars, will the landing systems work?" The expert assigns a probability of 60/65 to a positive answer. For the corresponding Venus question, the expert also assigns a probability of 60/65. (In general, these assignments need not be equal.) The expert has also stated that given a successful launch, information about the clairvoyant's report for the landing system for one destination will not influence his probability assignment for the other destination: the probability assignments to these events are conditionally independent. From this probability assignment model we can calculate

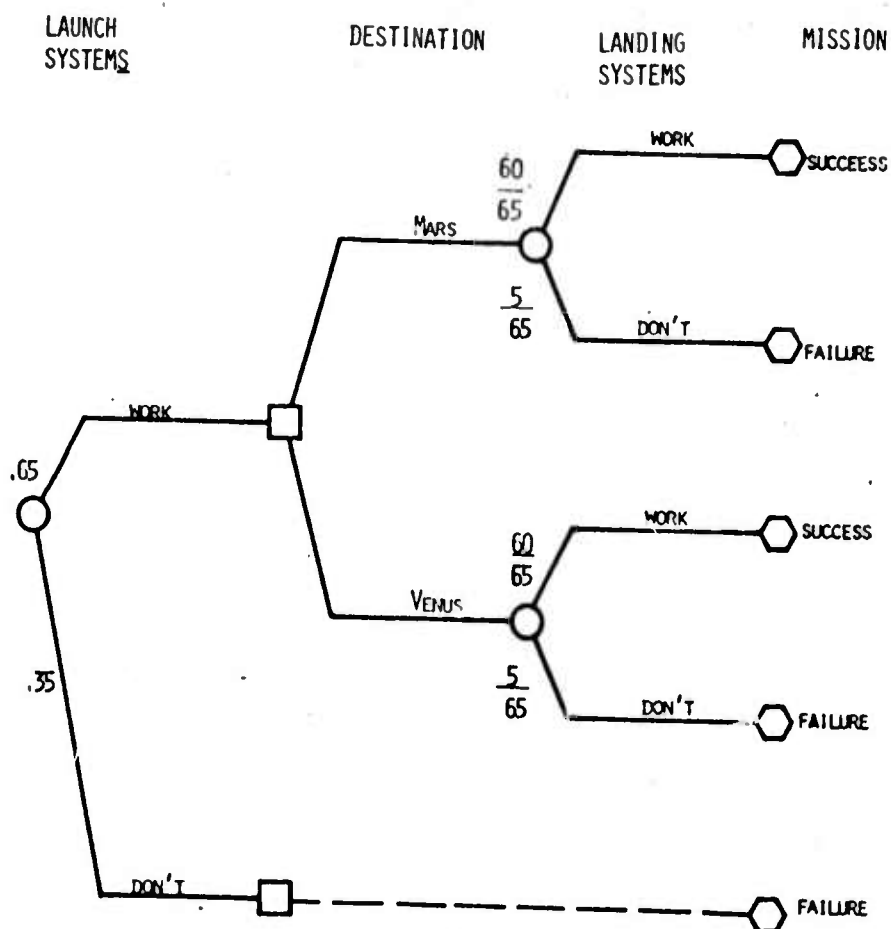


FIGURE 2.21: PROBABILITY ASSIGNMENT MODEL

the joint probability distribution for the clairvoyant's report of success or failure for the two possible missions (see Figure 2.19). For example, the probability that the clairvoyant will report success on both planets is:

$$0.65 \frac{60}{65} \frac{60}{65} = 0.554$$

Figure 2.21 is similar to what is commonly called the probability assignment tree, except that these trees usually do not include possible dependencies of probability assignments on decisions.

Further Implications of the Probability Assignment Model

The construction of a formal probability assignment model often raises new, interesting and useful informational questions. While the value of clairvoyance on the uncertainties appearing in the primary decision tree is an upper limit for the value of any corresponding information gathering program, often the most feasible information gathering programs are directly related to the uncertainties appearing in the probability assignment model. Thus, this model naturally leads to new and more practical information valuation questions.

For example, in the space mission problem we may be able to conduct exhaustive experiments on replicas of the launch system and make elaborate tests on the actual launch vehicle. The value of clairvoyance on the launch system alone is a straightforward calculation from the information we have built up.

The computation in Figure 2.22 shows that the value with clairvoyance on the launch system is 63 units. Subtracting the 56 unit value of the primary decision, we arrive at a 7-unit value of clairvoyance on the launch system only. Since most of the 9.84 value of complete clairvoyance can be derived from the more practical launch-system information, it would be best to start a realistic information gathering program with a study of the launch system.

Decision Dependent Clairvoyance

We may also wish to consider clairvoyance for only one of the decision alternatives. For example, suppose we engaged the clairvoyant to tell us only whether we will succeed or fail if we send the spacecraft to Mars. The primary decision tree of Figure 2.15 and the probability assignment model of Figure 2.21 give us all the information we need to construct the tree for the value with clairvoyance about Mars, shown in Figure 2.23. Once we get the clairvoyant's report on the success of the Mars mission, we must recalculate the probability of success on Venus because of the dependency in our joint probability assignment (which results from the common launch system). A report of a successful Mars mission results in a revised probability of 0.923 for success on Venus. A report of failure if we go to Mars revises the probability of failure on Venus to 0.885. Using these probabilities, we find the expected values shown in Figure 2.23. Contrary to our intuition, we find that a clairvoyant's report of a successful Mars mission indicates

CLAIRVOYANT REPORT ON LAUNCH SYSTEMS

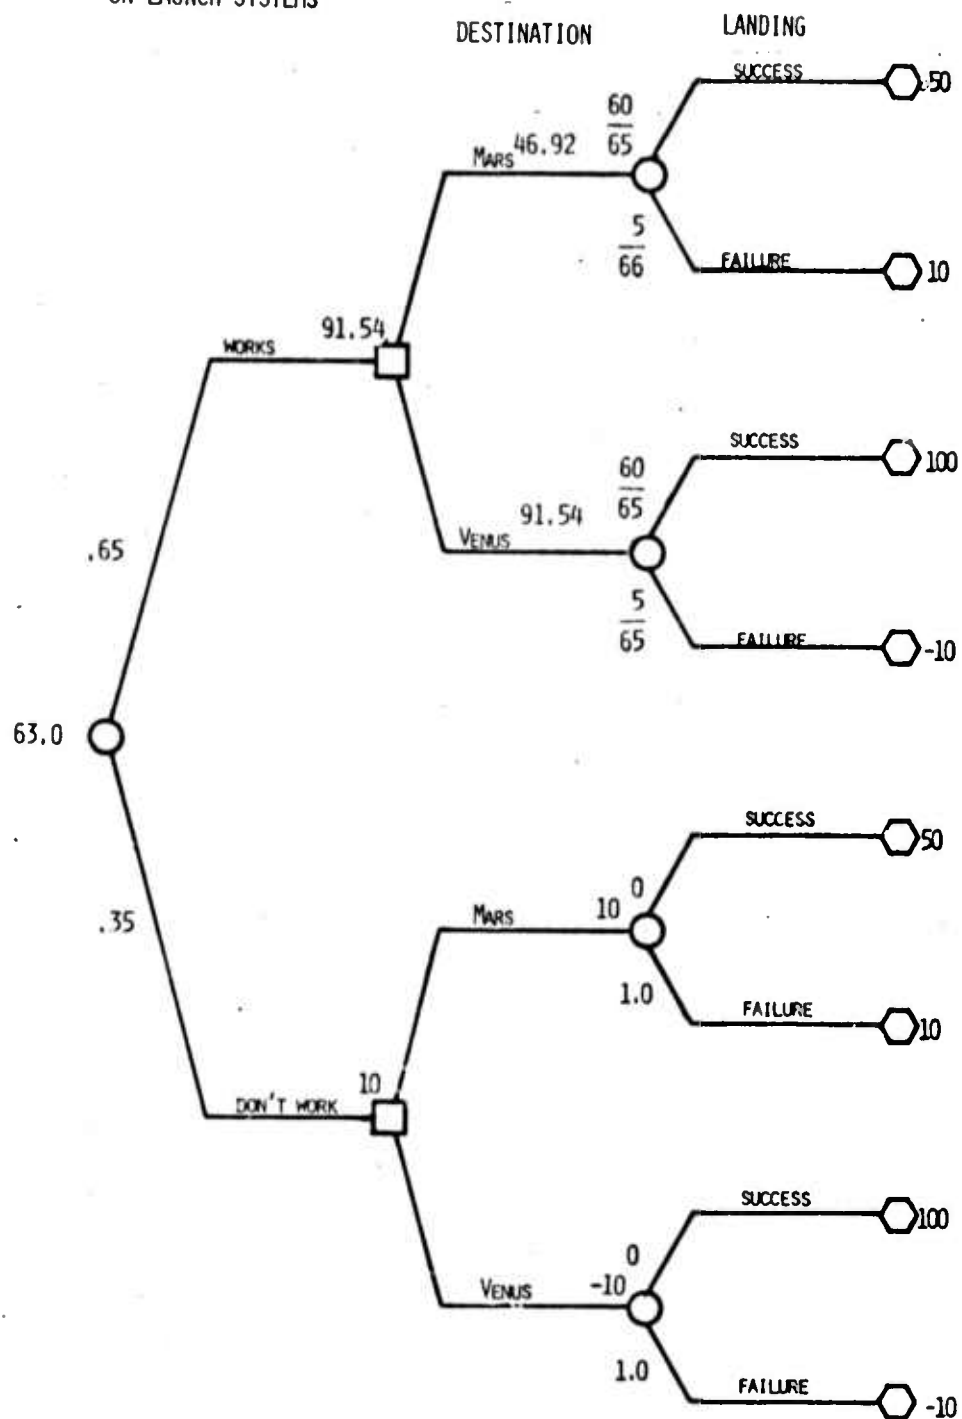


FIGURE 2.22: VALUE OF CLAIRVOYANCE ON LAUNCH SYSTEMS ONLY

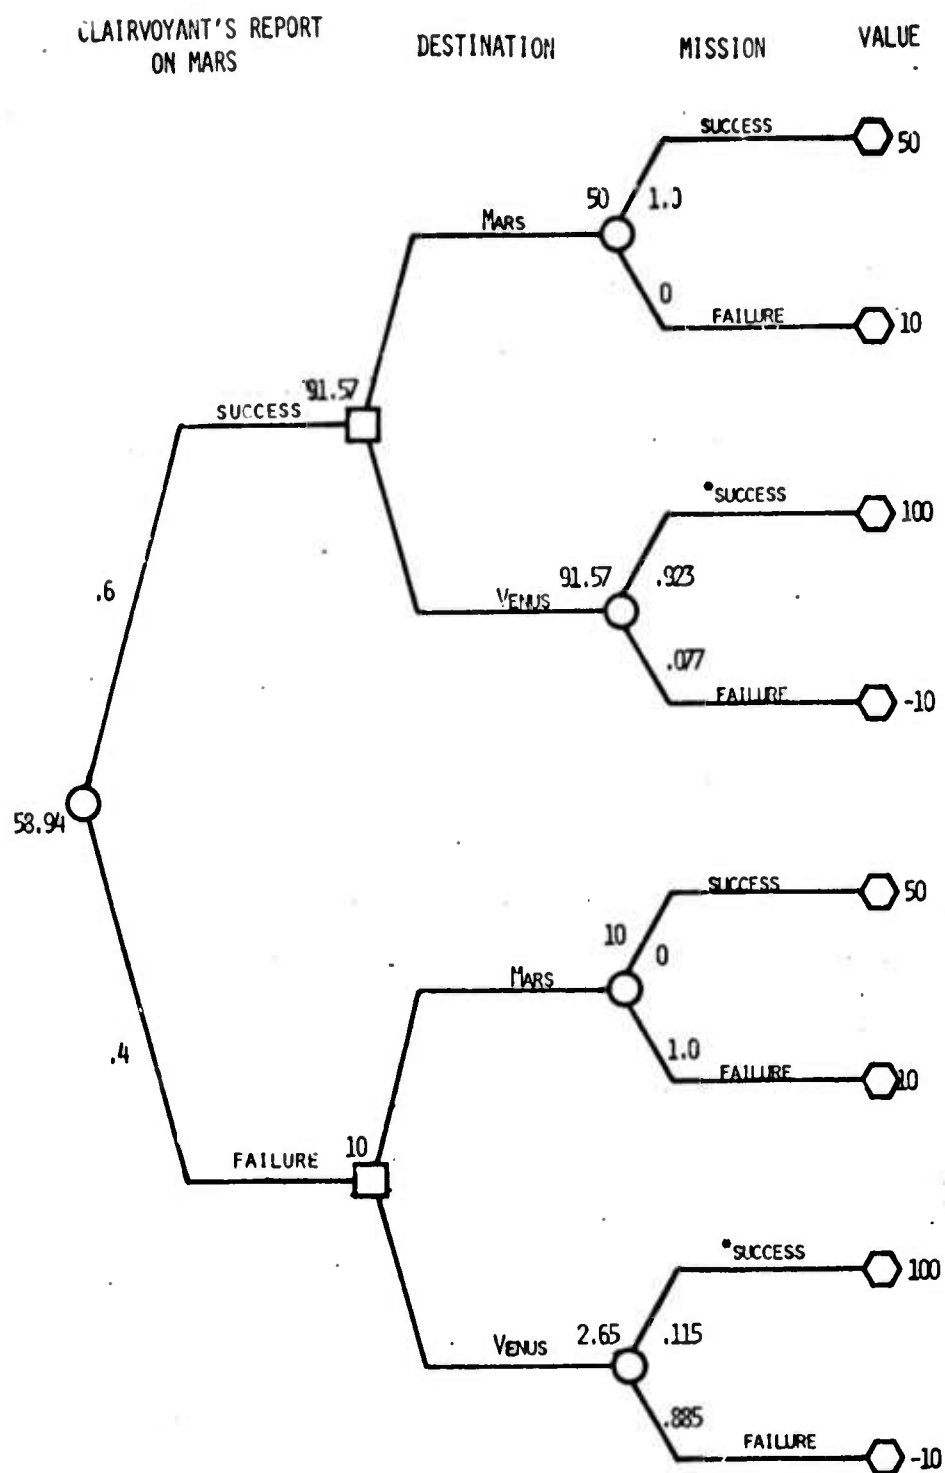


FIGURE 2.23: VALUE WITH CLAIRVOYANCE ON MARS ONLY

NOTE: SUCCESS PROBABILITIES ON VENUS ARE CHANGED BY MARS REPORT

that we should send the mission to Venus, and that a report that we will fail on Mars indicates that we should send the mission to Mars.

This is the phenomenon of decision-dependent information. Information about one aspect of a problem may have surprising implications for intuitively separate aspects of the problem due to dependencies in the probabilistic informational structure. In complex problems, a formal evaluation is the only way to determine the correct inferences and their implications.

In order to capture these effects, the analysis must not only represent the primary problem structure, (Figure 2.15), but it must also capture the informational structure in a formal model, (Figure 2.21). In many problems the informational model may provide the more natural and more productive focus for analysis. In the spacecraft example, we can derive the primary decision tree and all the informational trees from the single probability assignment model by adding only the decision-event chronology to apply to each case. This approach to analysis might provide a key to more effective computer aids to the model building process. (See Section 2.3.2, Analysis of Decision Trees.)

Unequal Decision-Dependent Probabilities

The calculations demonstrated above work equally well when the probabilities for success on Mars and Venus are not equal. We can demonstrate this fact by replacing the probabilities in Figure 2.21

with a "launch system working" probability of 0.75, a "Mars landing system works" probability of 70/75 and a "Venus landing system works" probability of 50/75. In the primary decision tree of Figure 2.15 this results in a "Mars success" probability of 0.7 and a "Venus success" probability of 0.5. The results with these new probability assignments, as well as the original ones, appear in Table 2.1 along with additional clairvoyance values for both landing systems and Venus mission only.

Conclusion

We have seen that a probability assessment model built on the concept of clairvoyance clarifies the interpretation and specification of probability assignments and precisely determines values of clairvoyance. It also adds precision to the subjective interpretation and assessment of probability. In problems where uncertainty plays a key role, emphasis on the construction of a formal informational model can clarify communication and lead to a more rapid and accurate solution.

Further research is needed to develop a precise and convenient notational system for dealing with probability assessment models. The existing inferential notational systems are too cumbersome and the existing graphical representations are incomplete. The general problem of specifying and translating between probability assessment trees and trees specifying the actual sequence of decisions and events is discussed in Section 2.3.2, Analysis of Decision Trees.

CLAIRVOYANCE ABOUT	EQUAL PROBABILITIES		UNEQUAL PROBABILITIES	
	Value with	Value of	Value with	Value of
Nothing	56	0	45	0
Everything	65.84	9.84	64.33	19.33
Launch System	63	7.0	50	5
Landing System	59.38	3.38	61	16
Mars Only	58.94	2.94	47.34	2.34
Venus Only	65.84	9.84	64.33	19.33

Table 2.1

SUMMARY OF SPACE MISSION EXAMPLE NUMERICAL
RESULTS WITH ORIGINAL (EQUAL) PROBABILITIES
AND NEW (UNEQUAL) PROBABILITIES

2.2.4 The Value of Sequential Information

Using decision analysis it is possible to calculate the value of one or more pieces of information--called "observables"--when a decision must be made in the face of uncertainty. This information has value because it can affect the decision and lead to a greater expected profit. However, the possibility of buying information sequentially presents the decision maker with a set of secondary decisions: which observables should he buy and in which order should he buy them? It is possible that knowing one observable affects not only the primary decision, but also the decision to buy additional information. In that case the value of knowing the first observable is greater than it would be if it affected only the primary decision. The prices of the observables affect the decision maker's willingness to buy additional information. For this reason the amount that the value of learning each observable is increased by the possibility of buying additional information depends on the prices of all the observables.

When the prices of all the observables can be added to determine the price of any combination of observables and when all the prices are known with certainty, we can formulate the general sequential--information problem in terms of a set of state variables (x_1, \dots, x_m) and a set of observables (y_1, \dots, y_n) with a corresponding set of observable prices $(K_{y_1}, \dots, K_{y_n})$. When an observable is equal to one of the state variables it represents perfect information.

However, by treating observables and state variables separately, we can also deal with imperfect information.

To solve for the value of information when all of the observables can be learned sequentially, we need to solve the decision tree shown in Figure 2.24. For a large decision problem it would be very difficult and tedious to generate this decision tree. However, the tree has a very repetitive structure that can be easily implemented as part of an automated decision aid for generating decision trees. Instead of the entire tree shown in Figure 2.24, the user specifies the decision tree that exists when information cannot be purchased sequentially and then asks the computer program to expand the tree to include sequential information.

The computer program starts the expanded decision tree with a decision node such as that shown at the left of Figure 2.24. The alternatives at this node are to buy any one of the specified set of observables or proceed to the basic decision tree without buying information. The last alternative leads directly to the basic decision tree specified by the user. The other alternatives lead to a chance node where the outcome of the selected observable is revealed.

After a chance node where one of the observables is revealed, the expanded decision tree contains a decision node where the alternatives are: to learn any of the specified set of observables that has not been learned previously, or to buy no further information.

Again, the last alternative leads directly to the basic decision tree specified by the user, except that this tree is now conditioned on the knowledge of one of the observables. The computer program continues to generate the expanded decision tree in this form until enough decision nodes are added to allow the decision maker to learn any subset of the specified observables in any order before proceeding to the primary decision problem.

By solving the decision tree in Figure 2.24 several times using different prices for each of the observables, it is possible for an automated decision aid to map out decision regions such as those shown in Figure 2.25. Figure 2.25 shows decision regions that might occur for two observables.

We cannot regard the value of learning one observable by itself as the maximum that we would be willing to pay for that piece of information. When it is possible to buy additional information sequentially, the value of an observable may increase. To determine an upper bound for realistic programs designed to gather sequential information, we need a decision aid that can generate and solve a decision tree like the one in Figure 2.24. Without this sort of aid, the problem structuring and computations are sufficiently difficult to discourage analysts from calculating the value of sequential information, even when the results might influence information-purchasing decisions.

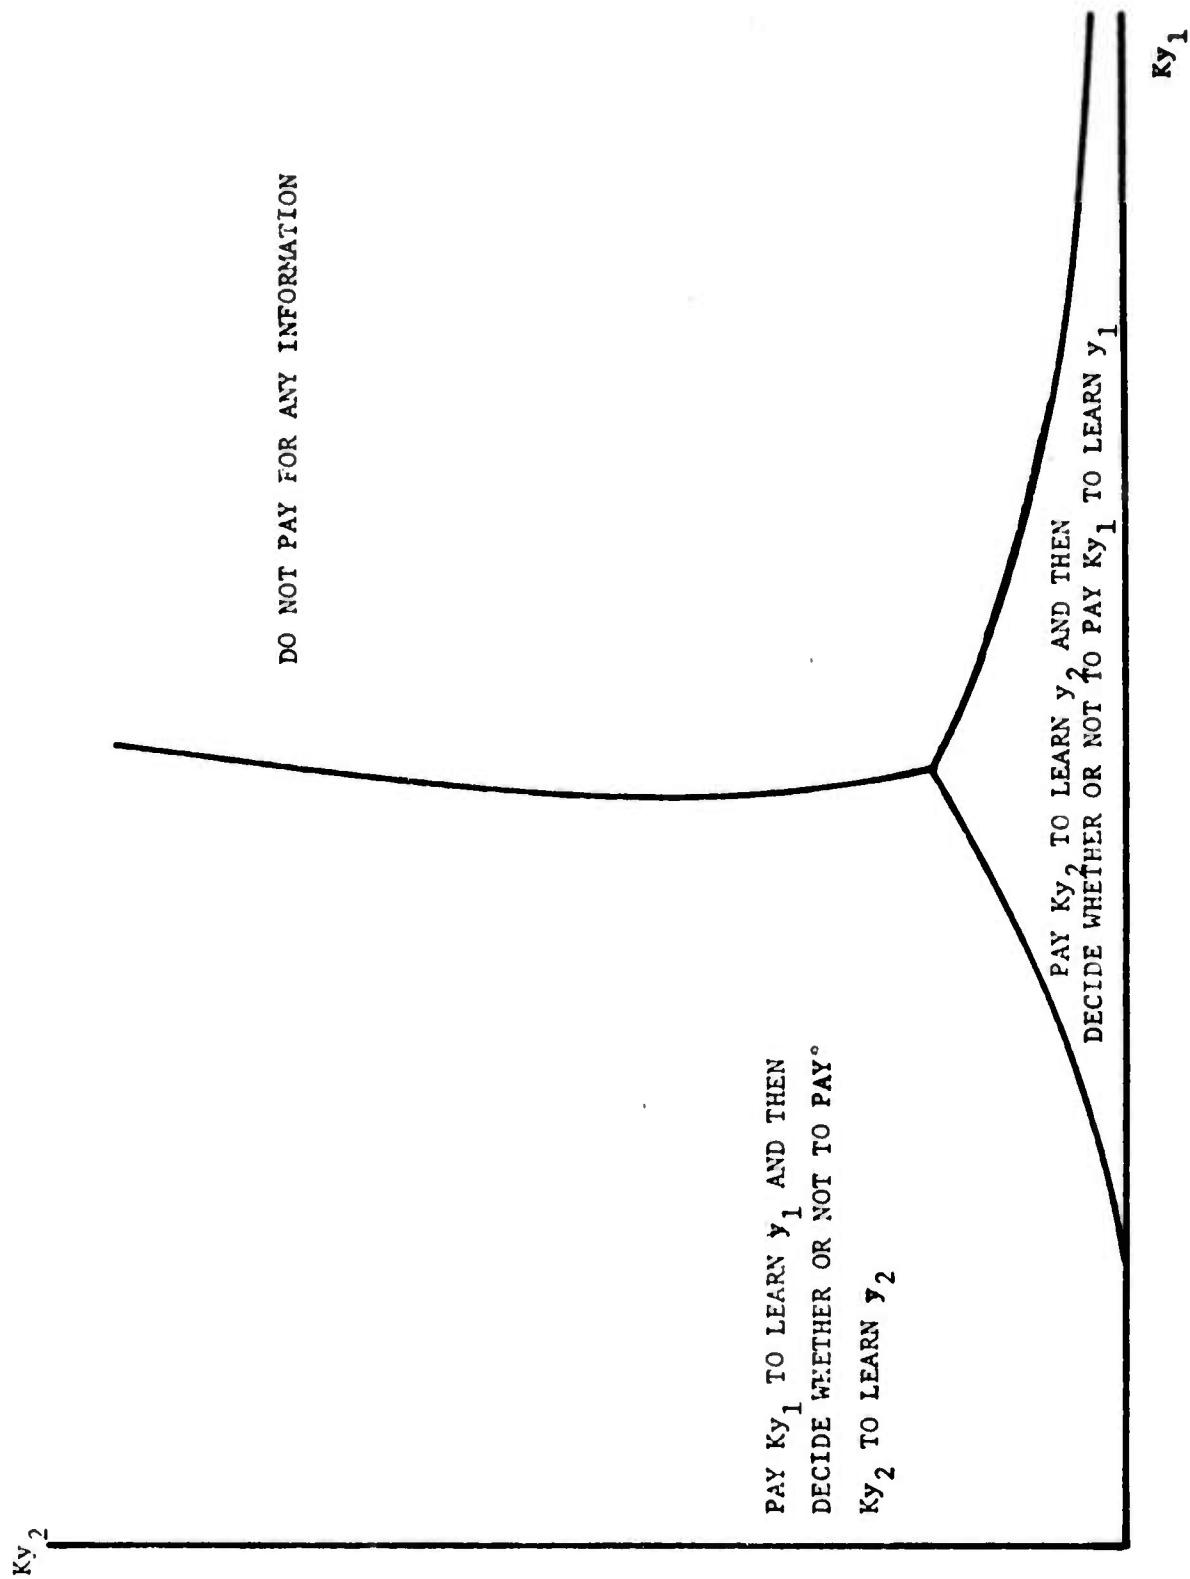


Figure 2.25: DECISION REGIONS FOR SEQUENTIAL INFORMATION

2.2.5 The Value of Flexibility

The notion that a good decision strategy is a flexible one has long been intuitively appreciated by decision makers. Nearly everyone is familiar with a story of a plan that went wrong because it failed to adjust for some unforeseen circumstance. Decision analysis has had little to say on the subject of flexibility. However, recent research on the concept of flexibility shows that this subject should be incorporated in a decision morphology.

Roughly speaking, something is flexible if it may be easily varied. However, in the context of decision making, ease of variation may be described by many different characteristics. Our point of view is that the flexibility of a given decision variable is determined by the nature of the choice set associated with that variable. The larger the choice set, the greater the decision flexibility. If the choice set consists of a single point element, in other words if the decision has already been committed, we say that the decision variable is inflexible.

A number of the classical micro-economic decision problems for which flexibility is a concern may be treated within this framework. Merkhofer [9] has shown that the problem of sizing a production facility can be so analyzed.* The decision strategy of committing something less than one's total resources so as to be prepared to meet unforeseen opportunities may also be expressed as a problem of maintaining flexibility.

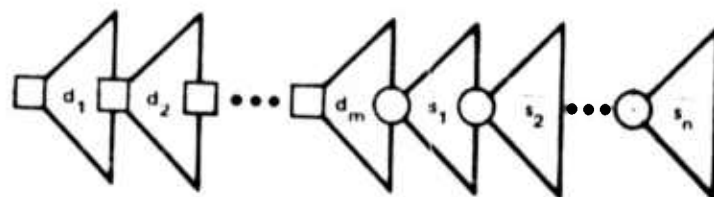
* For a discussion of this problem see Marschak and Nelson [8] and Baumol [1].

The value of flexibility is strongly dependent upon the information that might be received during the decision process. The more a decision maker expects to learn in the course of a decision, the more it pays to follow flexible decision strategies. Similarly, the more flexible one's decision strategy, the greater the value of information gathering. Thus, the concepts of value of information and value of flexibility become special cases of the more general concept of the value of information given flexibility.

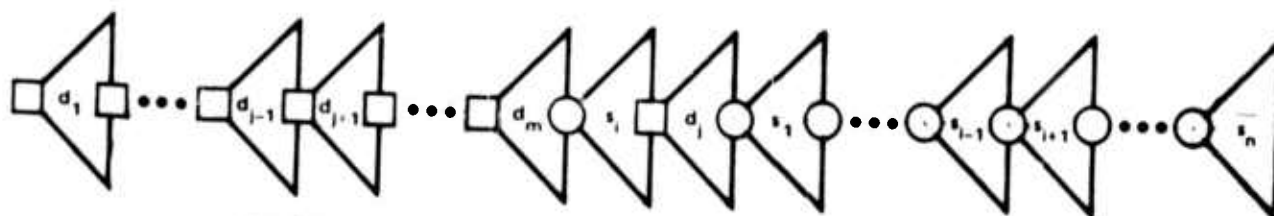
The value of information given flexibility measures the value to the decision maker, in economic units, of obtaining a given amount of information together with a given amount of decision flexibility. An upper limit to this quantity is the expected value of perfect information given perfect flexibility, $EV'IGPF$. Figure 2.26 illustrates the calculation of the $EV'IGPF$ in a decision tree.

Figure 2.26a shows a one stage decision problem. The decision maker must set a number of decision variables, denoted d_1, \dots, d_m . Subsequently, the outcomes of a number of random variables, s_1, \dots, s_n , become known. Once he has made his decision and the information concerning the random variable is revealed, the decision maker does not have the ability to go back and alter his decision settings. The structure of the decision tree in Figure 2.26a implies that our decision maker will receive no information prior to setting his decision variables.

Figure 2.26b illustrates the same components of the decision problem with perfect information on state variable s_1 given perfect flexibility on decision variable d_j . In this case the decision maker



(a) NO INFORMATION OR FLEXIBILITY



(b) PERFECT INFORMATION ON s_i GIVEN PERFECT FLEXIBILITY ON d_j

FIGURE 2.26 A PORTION OF A DECISION STRUCTURE

will learn the value of the i 'th state variable before he must set the j 'th decision variable. By calculating the maximum utility of the decision problems illustrated in both parts of Figure 2.26, the value of the information given flexibility can be obtained. For an expected value decision maker, the EVPIGPF will be the difference between the expected values associated with the two decision problems.

Thus, we see that calculation of the EVPIGPF involves the rearranging of decision and state variable nodes in the problem's decision tree. Therefore, the calculation of the value of flexibility, like the calculation of the value of information, is a tedious calculation.

Assistance in the form of a computerized system for restructuring decision trees would be useful for information and flexibility computations. With such an aid, the analyst would specify the potential information variables and those decision variables that could be set in response to that information. Tree restructuring would then be performed automatically and the new decision structure evaluated. The computed output would be the expected value to the decision maker of obtaining that combination of information and decision flexibility. This value would be extremely useful to the decision maker for evaluating various proposed information gathering and distribution systems.

2.2.6 Topics for Further Research on the Properties of Probabilistic Models

There are several other areas where further research is needed to understand the theoretical properties of probabilistic models so that these properties can be used as the basis for efficient, automated decision aids. In each case we need to develop ways of characterizing and visualizing the important elements of the problem in such a way that they can be expressed in a language suitable for automated model generation and analysis.

Recent research into joint time-risk preference has demonstrated that simplified procedures can be used to deal with complex decision problems involving the allocation and consumption of resources over long periods of time. Automated decision aids based on these procedures may, for the first time, make it practical for a decision analyst to construct and analyze large decision models incorporating complex preferences concerning the delay of consumption and variation in consumption patterns.

Another area where recent research has led to theoretical developments that may prove to be the basis of new automated decision aids is the analysis of competitive decisions, especially when the competition takes the form of bargaining or negotiations. There is still a long way to go before research on competitive decisions (game theory, social choice theory, bargaining theory) can form the basis of practical analytical procedures. However, we must consider the possibility of using such procedures as part of a decision morphology.

2.3 ANALYSIS OF DECISION MODELS

Once a decision model has been constructed and the necessary quantitative inputs (probabilities and values) have been elicited, a number of different analyses can be performed using the model in order to gain insight into the decision problem. The purpose of the following sections is to discuss at a conceptual level the various analytical procedures that can be applied to decision models, without dwelling on the current capability of computer systems to carry out the analyses. These conceptual tools, together with those discussed in Section 2.2 (Theoretical Properties of Probabilistic Models), will form the basis of computer programs that analyze decision models.

2.3.1 Sensitivity Analysis

One of the basic tools of present-day decision modeling is sensitivity analysis. By determining the sensitivity of one of the model parameters to changes in another model parameter, we can gain some insight into the relative importance of various components of the model. However, some of our recent research has shown that sensitivities may be a poor guide in determining, for example, the value of perfect information about an uncertain model parameter. Furthermore, we currently lack adequate procedures for calculating and visualizing joint sensitivities: the sensitivity of one model parameter to changes in several of the other model parameters.

There are a number of ways to measure the sensitivity of one parameter to another. If both parameters are known with certainty, then we can use a deterministic sensitivity analysis. The most common way to carry out this form of analysis is to vary one of the parameters over a specified range, and plot the corresponding values of the second parameter as a function of the first parameter. The resulting graph, showing the pairs of parameter values that fit the model, describes the global sensitivity of the parameters to each other. The analysis is global in the sense that the relationship between the two parameters is described over a given range. On the other hand, we could compute the local sensitivity of one parameter to another, at a given operating point, by determining the rate at which small changes in one parameter cause the other parameter to change. Local sensitivities are measured in terms of partial derivatives.

Another way to classify sensitivity analyses is by whether they

are open-loop or closed-loop sensitivities. A closed-loop sensitivity is one in which the decisions embedded in the model are reoptimized each time one of the parameters is changed. An open-loop sensitivity, which is usually easier to compute than a closed-loop sensitivity, is one in which the decisions in the model are fixed throughout the analysis even though changing one of the parameter values over its range might change the optimum decision.

For uncertain model parameters, a probabilistic sensitivity analysis is possible. In its most common form the probabilistic parameter is held constant at each of its possible outcome levels and the resultant effect on the outcome lottery is observed. However, there are a number of other possibilities. In general, a probabilistic sensitivity involves the modification of the probability distribution for one model parameter and the observation of the effect on the distribution for other model parameters. The procedures for probabilistic sensitivity analysis are not as well defined as those for deterministic sensitivity analysis because there are a variety of ways in which one can modify a probability distribution. For example, in one analysis it might be desirable to vary the variance of one distribution while holding its mean constant; in another analysis it might be preferable to vary algebraic coefficients in the expression for the probability distribution. As with deterministic sensitivity analyses, probabilistic sensitivity analyses can be either global or local, and either open-loop or closed-loop.

Further research is needed to determine the proper use of sensitivity analysis in the construction and analysis of decision models. While sensitivity analyses have a strong intuitive appeal, they can supply at best a rough idea of the implications of various decision models. We need a better understanding of the linkages between sensitivity analyses and other analytical measures such as the value of information and flexibility, and we need better guidelines for when it is appropriate to apply the various forms of sensitivity analyses.

2.3.2 Analysis of Decision Trees

Decision trees are one of the fundamental tools of decision analysis. However, a characteristic of decision trees is that they tend to grow very rapidly as the size of the model is expanded. At present it is not uncommon for decision analysts to utilize decision trees containing thousands, or even millions, of nodes. An analyst often finds that straightforward procedures for visualizing, modifying and analyzing small decision trees become far too time consuming when the size of the tree grows, even when automated decision aids are available.

To visualize a large decision tree, the analyst must often think in terms of simplified representations of the tree. The analyst constructs the decision tree by first building the simplified representation, and then using it to generate the entire tree structure. (See the discussion in Section 2.1.1 Algebraic and Graphical Languages for Decision Models.) Once the entire decision tree has been generated, the analyst may wish to check his logic by looking at various portions

of the decision tree. To do this he needs a procedure for rapidly selecting a particular location in the decision tree and then looking at the tree structure surrounding that location, preferably presented in graphical form. He may also wish to "fly over" the decision tree following a particular path through the decision tree or moving across the decision tree to investigate similar nodes.

During this process the analyst may find an incorrect portion of the decision tree indicating that the logic he used to construct the tree was faulty. At this point he might want to make a local change in the full-scale decision tree. However, since an error in one part of a large decision tree is often symptomatic of similar errors elsewhere in the tree, it would probably be better for the analyst to return to his simplified representation of the decision tree and correct the error at that level. Once the error has been corrected in the simplified version of the decision tree, the entire decision tree can be regenerated and the analyst can check to see that his changes have produced the desired results.

A capability that will undoubtedly be needed in future decision aiding systems is automatic processing of probabilities within a decision tree using Bayes' rule. This form of decision-tree processing can be implemented through the use of two decision trees for the same problem: one decision tree showing the sequence of events and decisions faced by the decision maker, and another decision tree (sometimes called a probability assessment tree) that specifies the same events and decisions in the order in which it is easiest to assess the necessary probabilities.

If the decision problem is complex, each of these decision trees might be accompanied by an appropriate simplified representation. In

addition to the procedures needed to translate from each simplified representation to the corresponding decision tree, general procedures based on Bayes' rule will be needed to translate from the probability assessment tree to the decision tree that specifies the sequence of decisions and events.

As the discussion in Section 2.2.3, The Value of Decision Dependent Clairvoyance, has indicated, it is difficult to apply Bayes' rule without carefully considering the structure of the decision problem. However, by including these considerations as part of an automatic system for Bayes' rule processing, it should be possible for an analyst to deal directly with a problem in the form in which it is easiest to assess probabilities (the probability assessment tree) and then translate automatically to a decision tree that specifies the events and decisions in the order in which they occur.

There are often several different ways to compute a particular result using a decision tree. For example, to determine the profit lottery associated with a decision tree, it is possible to work forward through the tree until every possible outcome has been determined together with its associated probability. Alternatively, it is possible to work backwards through a decision tree determining the profit lottery at each node in the tree until the profit lottery at the beginning of the tree has been determined. The structure of the tree, especially the amount of coalescence, determines the most efficient solution technique. It would be desirable for automated decision aids

to contain several different solution techniques together with algorithms for testing the structure of a decision tree and determining the most efficient analytical techniques.

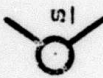

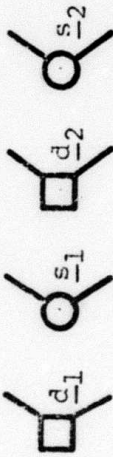
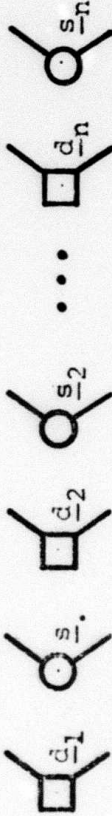
2.3.3 Approximate Methods of Probabilistic Processing

The purpose of this section is to explain how approximate methods can be used to determine the value of information and manage the growth of models. The need for model management to curtail exponential growth of model size is discussed in the Section 2.1.3. The approximate value of information calculations are based on deterministic sensitivities, and decisions to increase or decrease the size of a model are based on the approximate values of information.

Background

Decision problems may be grouped according to the number of stages they contain. We define a half-stage problem as one without decision variables. A single stage is any number of decision variables followed by any number of state variables. Multistage problems are characterized by decision variables separated by state variables. Figure 2.27 classifies typical problems. Any decision problem can be solved to any desired accuracy using trees; however, in many problems certain "smoothness" properties can be exploited to find answers as accurate as those from trees at a fraction of the computational cost. For one-half or one stage problems specific procedures can be programmed that are adequate for the majority of problems. Two stage procedures exist for a class of problems. Practical procedures do not exist for many stage problems.

Figure 2.27: CLASSIFICATION OF DECISION PROBLEMS

NUMBER OF STAGES	DIAGRAM	COMMENTS
1/2		Change of variables
1		Most common decision problems
2		Experimental version of 1 stage, pilot plant plus option for full scale plant
many		Dynamic programming problems, portfolio problems.

NOT REPRODUCIBLE

For simplicity we will treat the case of risk indifference in this section. However, T. Rice [10] has shown that only minor modifications are required to extend the framework to include risk aversion.

A value function, also known as a deterministic model or a profit function, can be expanded in a Taylor series about the mean of the state variables and the deterministic optimum setting of the decision variables. This approach is exact for a single stage problem that is quadratic in continuous state and decision variables and for discrete decision problems where the value function is a linear function of the state variables.

R. Howard has developed and documented this approach for change of variable problems in his course notes for EES 221 [4b] at Stanford University and for one stage problems in his paper Proximal Decision Analysis [4c]. His results are algebraic expressions for the mean and variance of the profit lottery and for the value of clairvoyance on the state variables. Means, variances, and covariances for the state variables, as well as the partial derivations of the profit with respect to the state and decision variables are required inputs. To automate Howard's methodology, the computer would be given the value or profit function and the moments of the state variables. Then the computer would automatically run sensitivities to estimate the required partial derivatives.

Rice has generalized Howard's approach. The general approach also starts with sensitivities and ends with the value of clairvoyance on each state variable. However, the middle step is eliminated, going

directly from sensitivities to value of clairvoyance without evaluating partial derivatives. For continuous problems, ones with continuous state and decision variables, the two methods are equivalent. The advantage of the direct method is that it also works for problems with discrete decision and/or state variables and with discontinuous value functions.

Computerizing the Single Stage Model

For a computer to find the value of clairvoyance for a state variable given only a single stage profit function and the moments (or distributions) of the state variables, two conditions must hold:

- (i) The state variables must be probabilistically independent of the decision variables.
- (ii) The value structure must be of the form $v(s,d)$: a deterministic model which assigns a single profit measure to each complete vector of state and decision variables.

Neither condition is restrictive. The Entrepreneur's Problem in Proximal Decision Analysis [4c] demonstrates how apparent violations of (i) can be rectified by reformulating the deterministic model. The model form $v(s,d)$ in (ii) is common to many discrete as well as continuous problems.

The basis for the computer program is the approximate value of information. The expected value of clairvoyance on the i 'th state variable is the prior expectation of the conditional expected value given s_i less the prior expected value:

$$\langle v_{c_i} | \xi \rangle = \langle \langle v | s_i, \xi \rangle | \xi \rangle - \langle v | \xi \rangle$$

Alternatively, we can focus on changes of decision by subtracting the prior expected value from the conditional expected value before

taking the second expectation:

$$\langle v_{c_1} | \xi \rangle = \langle \langle v | s_1, \xi \rangle - \langle v | \xi \rangle | \xi \rangle$$

The value of clairvoyance is the expected increase in value from making decisions after s_1 is revealed rather than before. We define the quantity:

$$\langle v | s_1, \xi \rangle - \langle v | \xi \rangle$$

as the stochastic compensation, stochastic because each term is an expected value and compensation because the decision is reoptimized to compensate for the departure of the i 'th state variable from its mean. Using the new terminology, the value of clairvoyance on the i 'th state variable is the expected stochastic compensation.

To approximate the value of clairvoyance deterministic compensation may be substituted for stochastic compensation.

Steps to Compute the Approximate Value of Information

To calculate the approximate value of information for the single stage problem, a computer must complete the following steps:

- 1) Accept the problem specification
- 2) Solve for the deterministic optimum decision
- 3) Perform deterministic open loop sensitivities
- 4) Perform deterministic closed loop sensitivities
- 5) Generate compensation functions
- 6) Compute the expected compensation

The procedure is remarkably robust. It works regardless of whether the state and decision variables are continuous or discrete. - However,

because of differences in optimization techniques the algorithms to achieve steps 1-6 are different for discrete and continuous decision variables.

Discrete Decision Variable

For a discrete decision variable steps 2-4 can be performed simultaneously. As illustrated in Figure 2.28 for two decision alternatives, the value v must be computed for each alternative at \bar{s} , and at $\bar{s} \pm \Delta s$. Then the open loop sensitivity is the curve with the highest value at \bar{s} , the one for d_1 , in the example. The closed loop sensitivity is the upper boundary of the curves. In Figure 2.28 it is the d_1 open loop sensitivity to the right of the crossover point and the d_2 open loop sensitivity to the left of the crossover point. The compensation is the difference between the closed and open loop sensitivities, plotted in Figure 2.29 for the example. The expected value of clairvoyance is found by numerically integrating the product of the compensation plot and the distribution on the state variable.

The extension of the procedure to many discrete alternatives is straightforward. There is one open loop sensitivity for each alternative. The closed loop sensitivity is the concave envelope of the open loop sensitivities. The compensation is the difference between the closed and open loop sensitivities as before.

Continuous Decision Variables

For continuous decision variables or for discrete decision variables with many alternatives, the computer will generate conventional sensitivities. Starting with the base case where all state variables are set to their means and all decision variables are set at the

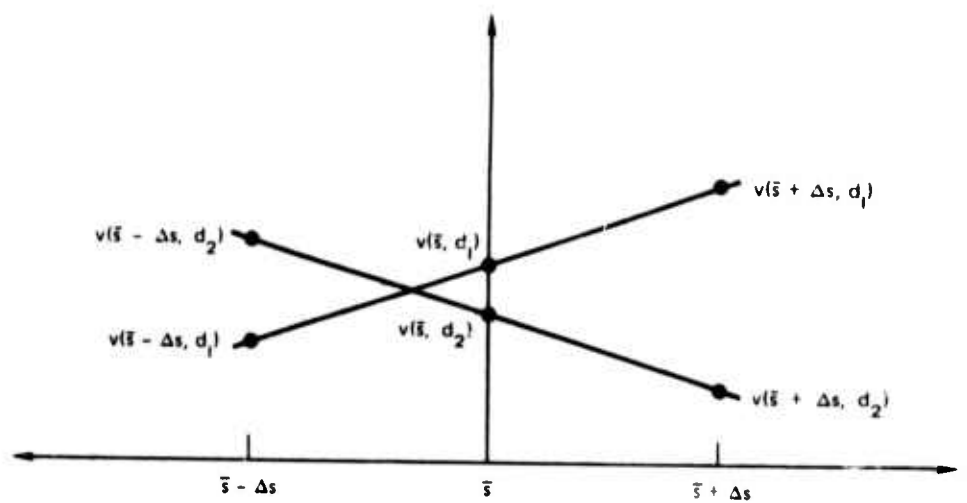


FIGURE 2.28 OPEN AND CLOSED LOOP SENSITIVITIES

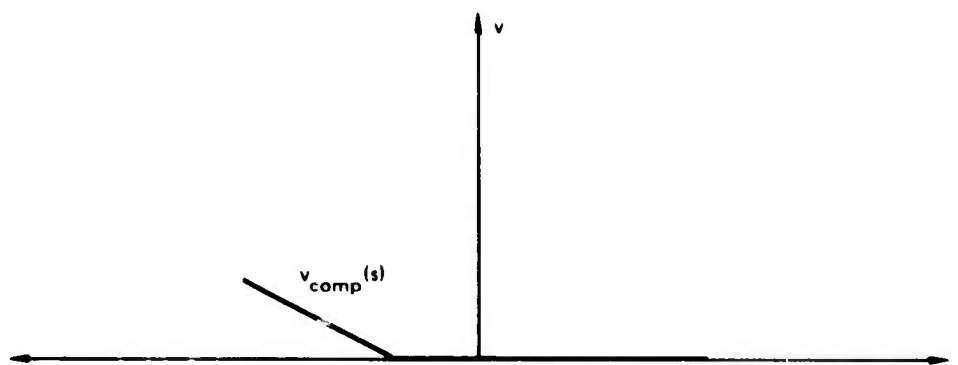


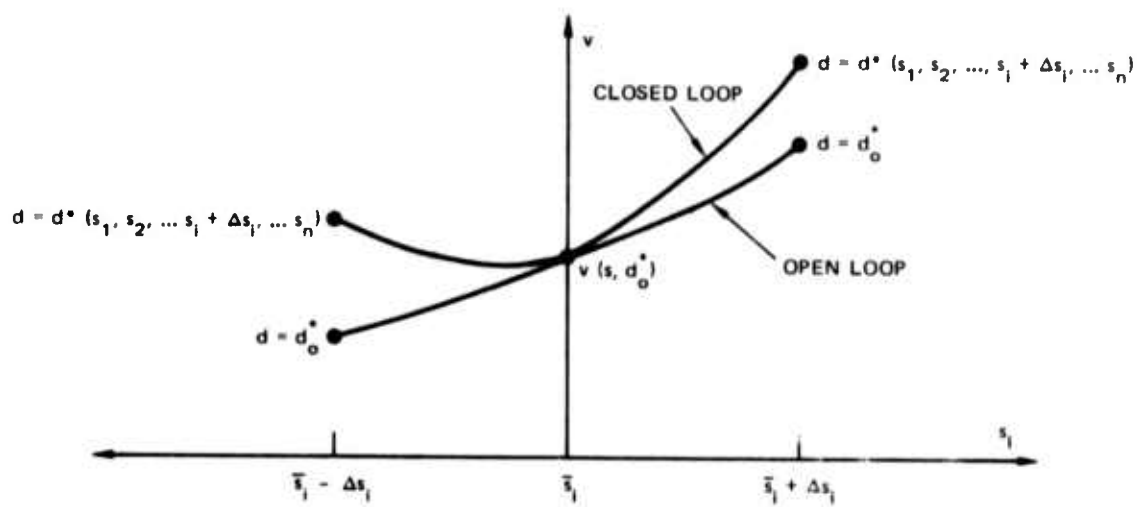
FIGURE 2.29 COMPENSATION FOR A DISCRETE DECISION VARIABLE

deterministic optimum, the computer successively adds and subtracts an increment to the i 'th state variable. As illustrated in Figure 2.30a, the decision is reoptimized at each point for the closed loop sensitivity and held at the deterministic optimum for the open loop sensitivity. The compensation plotted in Figure 2.30b is the difference between the open and closed loop sensitivities. To compute the expected compensation the computer will approximate the compensation with a quadratic or other curve form and perform numerical integration.

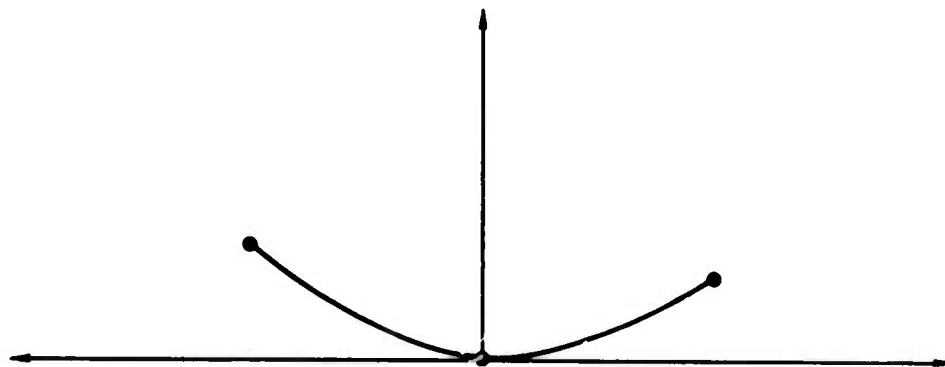
Two Stage Problems

Merkhofer [9] has shown that under certain conditions the two stage continuous decision problem illustrated in Figure 2.31 is solvable by these methods. As discussed in The Economics of Decision Making [10], the only important terms in the value function are the G matrix of second partial derivatives of value with respect to the i 'th state variable and the j 'th decision variable and the H matrix of second partial derivations of value with respect to decision variable i and decision variable j . These matrices may be positioned as shown in Figure 2.31. The value of compensation for this model is defined to be the difference between the open loop sensitivity of the value function to the state variables s_1 and the partially closed loop sensitivity in which the decision variables d_1 are held at their deterministic optimums but the decision variables d_2 are continuously optimized as s_1 is varied. The expected deterministic compensation is calculated to be:

$$\langle v_{\text{comp}} | \xi \rangle = - \frac{1}{2} \langle s_1' G_{12} H_{22}^{-1} G_{12}' s_1 | \xi \rangle$$

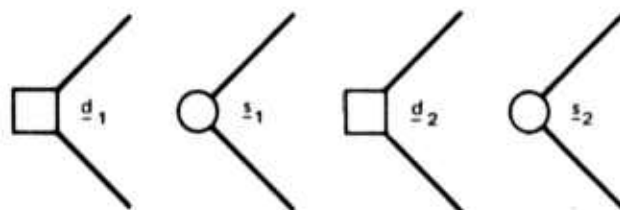


(a) OPEN AND CLOSED LOOP SENSITIVITIES



(b) COMPENSATION

FIGURE 2.30 COMPUTATION FOR CONTINUOUS DECISION VARIABLES



NOTE: $\left\{ s_2 \mid s_1, d_1, E \right\} = \left\{ s_2 \mid E \right\} ; \left\{ s_1 \mid d_1, E \right\} = \left\{ s_1 \mid E \right\}$

(a) THE SCHEMATIC TREE

$$\begin{aligned}
 v(2, \underline{d}) &= v(s_1, s_2, d_1, d_2) \\
 &= \underline{s}' \underline{G} \underline{d} + \underline{d}' \underline{H} \underline{d} \\
 &= \underline{s_1 s_2} \begin{bmatrix} \underline{G}_{11} & \underline{G}_{12} \\ \underline{G}_{21} & \underline{G}_{22} \end{bmatrix} \begin{bmatrix} d_1 \\ d_2 \end{bmatrix} + \underline{d_1 d_2} \begin{bmatrix} \underline{H}_{11} & \underline{H}_{12} \\ \underline{H}_{12} & \underline{H}_{22} \end{bmatrix} \begin{bmatrix} d_1 \\ d_2 \end{bmatrix}
 \end{aligned}$$

(b) VALUE FUNCTION

FIGURE 2.31 THE TWO STAGE QUADRATIC DECISION PROBLEM

The expected value of the optimal decision strategy will equal the expected deterministic compensation under either of the following conditions:

1. \underline{s}_2 is probabilistically independent of \underline{s}_1 .
2. G_{22} is a zero matrix and $\underline{s}_2 | \underline{s}_1, \epsilon \setminus \epsilon = 0$. (Non-observable state variables are deterministically independent of the flexible decision variables and the prior expectation is that the posterior mean of \underline{s}_2 will not be shifted by knowledge of \underline{s}_1 .)

These conditions are frequently satisfied in practice.

2.4 ELICITATION OF SUBJECTIVE INFORMATION

Automated decision aids have a potentially useful role in the elicitation of subjective information. However, the aids that have been developed to date in this area have been rudimentary at best.

The primary limitation of current aids is that they are designed to take a probabilistic "snapshot" of a well defined variable. They record probabilities and values without being sensitive to biases, conditionality, and shifts in the subject's opinion. Consequently they never recommend additional training to overcome biases, or additional structuring or modeling to allow conditioning on previously unstated variables, or additional encoding to replace data where the subject has simply changed his mind.

Although automated systems may eventually displace the human interviewer, initial work in this area will probably be to develop computerized aids to make the human interviewer more effective and efficient.

A second problem with current techniques is that they are insensitive to the reason the variable is being encoded. Presumably, within the context of a given analysis accuracy is more important for some variables than for others. The order of encoding and resulting conditionalities may affect the value of information calculations that can be performed in an analysis. In other words, the natural order of developing a model and then encoding the necessary probabilities applies to the development of automated aids as well as to the execution of a decision analysis project.

As a result we have spent most of our effort during the first part of this research project in areas other than elicitation. The following subsections contain some preliminary thoughts on the process of encoding subjective information. We intend to expand these ideas during the latter half of the current research effort.

2.4.1 Probability Encoding

The process of probability encoding has been the subject of considerable research, but very little of this effort has been directed toward the development of automated probability-encoding aids. At the present time almost all probability encoding takes the form of an interview or discussion between an analyst and one or more individuals who supply the necessary information. The analyst may have the assistance of other individuals during the encoding process, and he may use some very simple visual aids to help the subjects conceptualize their probability estimates. However, with the exception of a few unsophisticated computer programs, automated decision aids for probability encoding are not available for either the analyst or the individuals supplying the estimates.

One question for which we do not currently have a satisfactory answer is whether the analyst can be completely replaced by an automated decision aid during the encoding process. Regardless of how sophisticated the aid may be, it will not have all of the information available to an analyst. Both an analyst and a decision aid will have numerical information supplied by the subject, but only an analyst will be able to detect the visual and

emotional cues that indicate that the subject is becoming uneasy or confused during the encoding process. The question of whether decision aids should be designed to replace the decision analyst's role in probability encoding or simply assist an analyst with this task will probably not be resolved until we have implemented various automated probability encoding aids and have gained some experience with their use. We may find that the analyst's role in probability encoding can be performed by an automated decision aid after the subject has had enough training to avoid some of the basic problems associated with quantifying subjective information. However, an analyst may be required the first few times that a subject is faced with the encoding process.

There is always a choice between encoding the uncertainty in an important variable and modeling the problem further. At one extreme, it is conceivable that the final outcome of a project could be encoded directly, thus bypassing the need for examining the underlying variables of the problem. Generally, however, the distribution for the final outcome is more easily reached, and provides more confidence, if a model is constructed that relates the final outcome to other variables. The modeling effort tends to be most effective and economical if it starts with a simplified model that is successively refined. At some point during this process the model will become sufficiently complex to shift the balance between modeling and encoding. At this point the difficulties associated with expanding the model further outweigh the simplifications of the encoding process that would result from such an

expansion. In the discussion that follows it will be assumed that the trade-off between modeling and encoding has already been considered and the decision has been made to encode the necessary information.

Individual Probability Encoding

The majority of the research to date on probability encoding has dealt with the question of eliciting information from an individual subject. This form of encoding must overcome the cognitive and motivational biases that distort the subjective information supplied by an individual. Cognitive biases are a systematic distortion of an individual's subjective estimates caused by the way he thinks about uncertainty. For example, a response may be biased toward the most recent piece of information simply because that information is easiest to recall. Motivational biases are distortions in an individual's subjective judgment caused by his perceived system of personal reward for various responses. For example, an individual may want to bias his response because he perceives that his subsequent performance will be evaluated by a comparison of his response to the actual outcome.

Techniques for eliciting information may be classified according to the type of questions asked. Probability methods require the subject to assess the probability associated with a particular outcome of a state variable. Value methods require the subject to assign the outcome which corresponds to a particular probability. Mixed probability-value methods require the subject to assess

values on both the outcome and probability scales; the subject essentially describes points on the cumulative distribution of an uncertain quantity. These encoding methods can be used in either a direct or an indirect response mode. In the direct mode, the subject is asked questions that require numbers as answers. The answers can be given in the form of probabilities (or equivalently in the form of odds) or values. In the indirect response mode, the subject is asked to choose between two or more uncertain lotteries. The lotteries are adjusted until he is indifferent, and the point at which he is indifferent is translated into a probability or value assignment. The indirect response mode is typically used with a reference process, where the subject is asked to compare some aspect of an uncertain quantity to a reference process such as the toss of a fair coin or the spin of a wheel of fortune.

The most appropriate encoding method depends on the type of uncertainty being assessed. If there are only a few possible outcomes for an uncertain quantity, a method that requires the subject to divide up the range of possible outcomes into a number of intervals may not be appropriate. An automated decision aid should be capable of using any one of several encoding methods to elicit subjective information. Its choice of a particular encoding method should be based on the characteristics of the uncertain quantity, its importance to the modeling effort, and the personal preferences of the subject supplying the information.

The encoding process consists of five distinct phases: the motivational phase, structuring phase, conditioning phase, encoding phase, and the verification phase. Each may require different types of decision aids. In the motivational phase, it is necessary to explore the subject's motivational biases and attempt to eliminate or compensate for them. In the structuring or definition phase, the subject and analyst (or decision aid) must reach an agreement on the exact definition of the uncertain quantity being considered. The conditioning phase is directed toward finding out how the subject goes about making his probability assignments and heading off any biases that might surface during the encoding process. Once the uncertainty has been well defined and the subject's cognitive and motivational biases have been explored, the process enters the encoding phase and utilizes one of the encoding methods discussed above. In the verification phase the encoded information is subjected to a number of consistency checks to see if it truly represents the subject's beliefs.

There is little value in designing a decision aid that simply asks the subject for probability estimates. To do an adequate job of probability encoding the decision aid must be capable of working interactively with the subject to understand how he goes about making his probability assignments. This information can then be used by the decision aid to either train the subject to overcome the various biases that can distort his subjective estimates, or to devise means for compensating for the biases.

Encoding Probabilistic Dependencies

Although considerable research has been devoted to the question of encoding probabilities, very little work has been done on the question of encoding probabilistic dependencies. Decision analysts have recognized that the degree of difficulty associated with the encoding process is greatly increased when the probabilities being encoded depend on other random variables or decisions.

For example, consider a model in which there are four random variables each of which has three possible outcomes. If the random variables are all independent we can specify the uncertainty in the model with 12 probabilities. On the other hand, if the four random variables are all dependent, it will be necessary to assess a minimum of 81 probabilities.

In large probability models, the number of probability assessments required to specify the model can grow exponentially with the number of random variables in the model. Therefore, unless steps are taken to overcome the difficulties associated with encoding probabilistic dependencies, the number of probability assessments required can become prohibitively large. (Section 2.2.1, Probabilistic Dependence, and Section 2.4.2, Coalescence, discuss some of the conceptual tools that have been developed for describing and limiting probabilistic dependencies.)

In some situations it is easiest for a subject to assess the uncertainties associated with dependent random variables as joint probabilities, while in other situations the subject may find it easier to assess the same information in terms of conditional prob-

abilities. A joint probability distribution specifies the probability of occurrence of each possible combination of values of the random variables. Conditional probability distributions specify the probability that some subset of the dependent random variables have certain values given the values of the remaining random variables. In the example with four random variables, each of which has three possible outcomes, there are 81 joint probabilities, one for every combination of outcomes. In the same example, there are many possible conditional probabilities, of which 120 ($3 + 3^2 + 3^3 + 3^4$) would be required to specify all of the uncertainties in the problem. In this example, the number of conditional probabilities needed to specify the uncertainty exceeds the number of joint probabilities. However, in many large decision models containing some form of independence among the random variables, the number of assessed probabilities can be minimized by eliciting conditional rather than joint probabilities.

It is often difficult to conceptualize the probabilities associated with dependent random variables, especially when the random variables represent sequences of events that can occur in any order. Subjects usually think of sequences of events in terms of conditional probabilities, but when the order in which the events occur is not specified, it is not clear how the subject should assess the conditional probabilities.

Another problem with assessing conditional probabilities is that the subject may anchor on his probability estimate given one set of conditioning events, and then shift his probability estimates

by an insufficient amount when the condition events are changed. Before automated aids can be developed in this area, further research is needed to expand our understanding of these problems.

Encoding Rare Events

Encoding the probability of rare events--events which have a very small probability of occurring--is difficult for two reasons. The subject is being asked to assess the probability or odds of an event with which he has, by definition, little experience.

In addition, the subject may have to distinguish among very small probabilities or odds (for example, a probability of one in one thousand as compared to a probability of one in ten thousand).

Procedures have been proposed for overcoming these difficulties. One such procedure is to relate the rare event to other uncertain events with which the subject is more familiar, and to describe the event in terms of its component parts or as the result of a sequence of other events. The purpose of this procedure is threefold: to gather together all of the information relevant to the event which may be within the subject's command, to see if the event can be re-defined in terms which no longer involve small probabilities, and to produce lists of related events that can be used as reference points in getting the subject to make relative statements about their uncertainty.

The subject is then asked to make relative judgments about the occurrence of a number of different events before dealing with very small probability numbers. Subjects can often say whether the probability of one rare event is greater or less than the probability

of another rare event, even though they would find it difficult to assign a numerical value to either probability. By doing so the subject can bound the probability associated with the rare event and then narrow the bound as needed for the modeling effort.

Since subjects often have difficulty thinking about the probabilities or odds associated with rare events, decision aids might be used to help them visualize the problem. A decision aid could be used to generate a reference event against which the unknown probability could be compared. For example, one chance in ten thousand could be visualized in terms of a grid displayed on a computer screen with one hundred divisions on each side. The subject could be asked to compare the probability of the rare event with the probability that he could correctly pick a square from the grid that someone else had chosen without his knowledge. Decision aids could also be used to show the relationship between a rare event and other related events that might lead to its occurrence. However, before we can begin to design decision aids for encoding rare events, we need a much better understanding of the process used by individuals to conceptualize and express small probabilities or odds.

Group Encoding

Another area of probability encoding about which we need a better understanding before designing decision aids is group encoding. Group encoding refers to the process by which the judgment of a group of people is quantified. There are many decision problems for which more than one expert is available to supply subjective

estimates of uncertain quantities, and there can be considerable differences of opinion among the various experts. Group encoding procedures are used to aggregate the estimates of several experts into one probability distribution that can be used as the input to a decision analysis.

Group encoding is currently carried out with a group of subjects, whose judgments may have been individually encoded prior to the group encoding session, and an analyst who monitors the flow of information among the subjects to see that all the relevant opinions on the subject are discussed and incorporated into the group consensus. Another individual may have the responsibility to review the group consensus and dissenting opinions, and decide on the final probability distribution that should be used in the analysis. This individual is typically the head of the organization that has employed the services of the group of experts.

Much of the research associated with group probability encoding has been concerned with the way in which information should be exchanged among the individuals supplying the subjective estimates. At one extreme, it has been suggested that the participants in the group encoding process remain anonymous, and that they send each other their subjective estimates but not the logic behind the estimates. This procedure has been called the Delphi technique; it is designed to reduce the influence of group members with dominating personalities and encourage each member of the group

to reach his own opinion. Alternative procedures allow the participants to meet face-to-face and/or exchange information freely on an unlimited basis.

Automated decision aids are potentially very valuable for group encoding since they can act as a communication medium, as well as an elicitation device. If it is desirable for the subjects to remain anonymous, then an automated aid could be used to elicit the necessary information from each individual, and then rapidly aggregate and present the information to everyone in the group. The format of the group encoding process could be much more controlled if the communication environment were an automated decision aid instead of a conference table. Further research in this area should deal with the various ways a decision aid could aggregate and display information useful in helping the group reach a consensus.

2.4.2 Encoding Values and Preferences

The development of automated decision aids for encoding values and preferences may have to wait until the theoretical basis for alternative encoding procedures has been established. A number of different encoding procedures have been used in recent years, some of which are inconsistent with certain formulations of utility theory.

For example, one encoding procedure requires the subject to indicate the value of a particular outcome (along a particular dimension) by deciding where the outcome stands on a scale of 100, where 0 corresponds to the worst possible outcome and 100 corresponds

to the best possible outcome. In this case, if the subject picks a score of 50, he is saying that the incremental value of this outcome relative to the worst possible outcome is about $1/2$ of the incremental value of the best possible outcome.

However, this approach is inconsistent with the formulation of utility theory upon which much of the decision analysis is based. One of the implications of utility theory is that when an individual's preferences have been described on a particular utility scale, any monotonically-increasing transformation of that scale will result in an equally valid utility scale describing the same preferences. This means that the statement that an individual likes one outcome half as much as another can be transformed to a different scale in which the individual likes the first outcome $1/3$ as much as the second. Which scale was the subject thinking about when he said that he liked one outcome half as much as the other? The subject would probably have a great deal of difficulty answering this question since his preferences are being measured on an arbitrary utility scale that has no physical interpretation.

Once the theoretical difficulties have been resolved, we can develop automated aids for encoding values and preferences that are both practical and theoretically sound. These aids could be especially useful for problems which require complex value judgments. For example, it is frequently necessary to elicit preferences over multi-dimensional outcomes. Aids for multi-dimensional value encoding could be designed to help an individual

decompose a difficult value assessment into several value dimensions (cost, time to completion, loss of life, political effects, etc.) that are relatively easy to think about. Preferences could be elicited individually on each value dimension and then aggregated to determine multi-dimensional value assessment. As an additional aid, trade-off boundaries among the various value measures could be graphically illustrated.

Further research on this topic is important since values and preferences are central issues in most public sector decision analysis. Initial efforts at automation will be to build support systems for a human interviewer rather than to fully automate encoding techniques.

2.4.3 Encoding Time and Risk Preference

Encoding risk preference means determining an individual's attitude toward uncertain situations in which there is a chance of outcomes with very serious consequences. When faced with a decision involving his well-being or major portions of his assets, a decision maker may be unwilling to "play the averages". Decision analysts use the concepts of risk aversion and utility functions to deal with this phenomena.

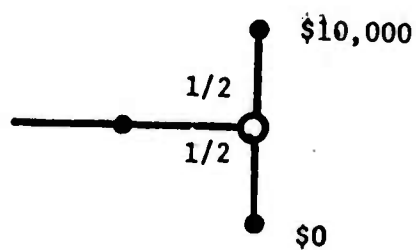
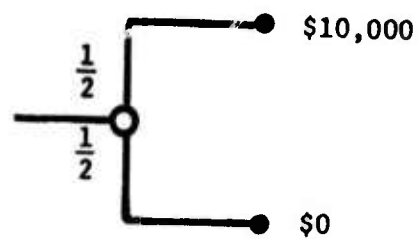
Most techniques for encoding an individual's attitude toward risk are designed to generate a utility function which specifies risk attitude over a range of possible outcomes. An encoding process, usually taking the form of an interview between the analyst and decision maker, is used to determine the utility function. The decision maker is asked to choose among a set of reference lotteries.

His choices are then analyzed to derive the utility function.

In the process of conceptualizing and expressing his attitude toward risk, a subject may unconsciously distort his beliefs in much the same way that he can unconsciously bias subjective probability estimates (see Section 2.4.1, Individual Probability Encoding). For example, most subjects exhibit the "zero illusion". This means that individuals tend to base their thinking about risky situations on what they perceive their current level of assets to be even though they may not know their current level of assets. Procedures are needed for overcoming this sort of distortion in an individual's assessment of risk preference.

The assessment of a decision maker's attitude toward risk becomes rather complex when the possible outcomes are distributed over time. It can be shown that an individual's attitude toward uncertain situations depends not only on the values of various outcomes and the time in which they might occur, but also on the time at which the uncertainties associated with the outcomes are resolved.

For example, consider the two lotteries shown in Figure 2.32. In both lotteries there is a 50% chance of winning \$10,000 and a 50% chance of winning nothing. Furthermore, in both lotteries the \$10,000 is received one year from the present time. However, in the first lottery the uncertainty is resolved immediately, while in the second lottery the lottery is not resolved for a year. When asked to choose between two lotteries, most individuals



$t = 0$ $t = 1$

→ time

Figure 2.32: AN EXAMPLE OF JOINT TIME RISK

will choose the first because immediate resolution of the uncertainty associated with the lottery allows them to plan their consumption over time.

Procedures are needed to determine the relative value of lotteries such as those shown in Figure 2.32. Recent theoretical research in time-risk preference has indicated that, for most problems, an individual's attitude toward both time and risk preference can be assessed in terms of a few simple parameters. Further research is needed to develop these procedures into a set of concepts that can form the basis for an automated decision aid designed to encode time-risk preference.

2.4.4 Generation and Elicitation of Alternatives

A good alternative can often be more valuable than the careful selection of a course of action from a set of poor alternatives. Executives faced with a difficult decision often spend a good portion of their time searching for a better alternative. However, very little research has been done to develop methods for helping the decision maker generate new alternatives and incorporate them into an analysis of the decision problem.

Determining the value of information and flexibility can be viewed as one way to generate new alternatives. In this case, the new alternative is the possibility of buying information before making the primary decision. However, while new information can cause a decision maker to pick a different alternative, it does not offer him any fundamentally different alternatives when he faces the primary decision.

It may be possible to elicit new alternatives from a decision maker by having him conceptualize possible combinations of existing alternatives and looking for hedging strategies that combine the desirable features of two or more existing alternatives. For example, an executive trying to decide whether or not to undertake a risky, but potentially very profitable, investment starts with two alternatives: to invest or not invest. It may be possible to elicit intermediate alternatives by having the subject consider various forms of risk sharing, which have the effect of reducing both the potential losses and the potential gains associated with the investment.

We need to know a great deal more about the generation and elicitation of alternatives before we can start designing decision aids to accomplish these tasks.

FIGURE 2.3A SIMPLE MODEL

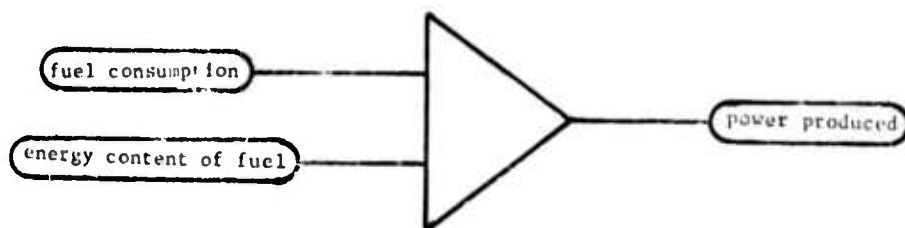


FIGURE 2.3B EXPANDED MODEL

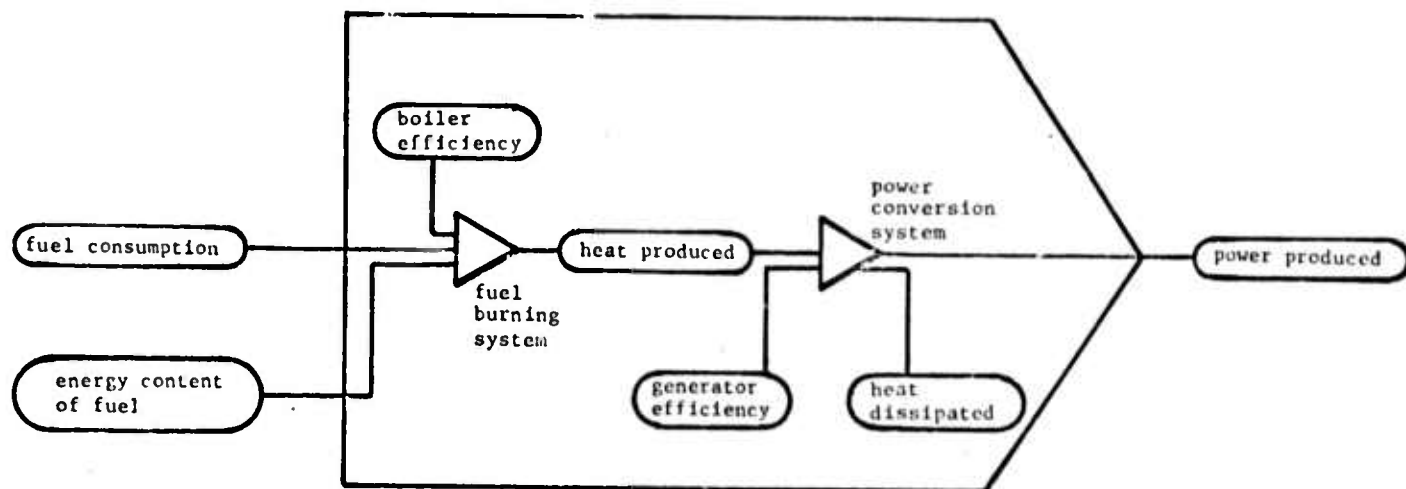


FIGURE 2.3 MODELS OF A COAL BURNING POWER PLANT

SECTION 3

AN ILLUSTRATIVE IMPLEMENTATION OF THE MORPHOLOGY: A COMPUTERIZED GRAPHICAL-ALGEBRAIC TECHNIQUE FOR MODEL BUILDING

In this section we shall discuss some preliminary results concerning the application of the decision morphology to the development of a system of computer aids for decision analysis. The first part of this section presents some requirements for a successful computerized system for decision making. The latter parts of this section describe how the graphical language discussed in Sections 2.1.1 and 2.3.2 might be used for model building and now this technique might be usefully implemented as part of a computerized system for decision analysis.

3.1 CHARACTERISTICS OF A SUCCESSFUL SYSTEM OF COMPUTER AIDS FOR DECISION ANALYSIS

Recent advances in computer hardware and software technology, including the development of visual aids and the introduction of time sharing systems, have tremendously increased the potential scope of computer analysis. User-machine interaction has progressed to the point where we now find ourselves in an era in which the computer's powerful processing and data-handling capabilities may be used to extend greatly human reasoning ability and insight.

Several attempts have been made to develop computer programs to be used as aids for decision modeling.* Certain deficiencies in these programs, however, have tended to limit their usefulness. One of the most significant limitations of existing systems is a lack of flexibility. Designing computer aids requires striking a compromise between the needs of the user and the difficulty of implementing the procedure on the computer. Too often the compromise is at the expense of the user.

We believe that we will ultimately achieve a better design if we initially ignore the constraints of existing computer systems, identifying first the basic requirements that the system must possess to be successful. Some of these requirements are essential to any computerized system that is designed to support modeling activity. If the system is designed specifically to aid decision making, it is desirable that it possess certain additional capabilities.

3.1.1 System Language Flexibility

As stated above, flexibility is an important characteristic for any system which is to be used in modeling. Of main concern for supplying flexibility is the adequacy of the system language for user-machine communication. The system language must permit the user to become directly involved with his problem rather than

*See, for example, Dahl and Nygaard [2], Groner et. al. [3], IBM [5], IBM [6], Kiviat et.al. [7], Zamora and Leaf [13].

with the computer. Therefore, the language of communication should be so natural that the user can devote his main attention to monitoring the status of his model. Graphics capability is an important aspect for many users, not only because plots are a natural means for representing data, but also because the relationships between model entities and operators are much easier to grasp if presented visually. The readability and formatting capabilities of a graphics display allow fast, selective scanning of large amounts of data. The user can switch quickly among various blocks of information to analyze diverse aspects of modeling.

The command language must be concise, convenient, and efficient in accessing all system facilities. It must be versatile, including, for example, the ability to stack commands before execution, to cancel a list of commands, and to define macro commands composed of other commands. A facility is also needed to allow extensions to the command language to provide for problem dependent operation.

It should be possible to tailor the system language to fit the class of models being constructed. When appropriate, the language should provide special purpose statements which allow direct and simple specification of all basic model characteristics. As much of the definition of the model as possible should be done for the user so he need supply only particulars and can concentrate on the unique parts of his system. In addition, the language should provide programmed routines for standard methods of defining

common characteristics of the model. Such routines for model definition are essential for large scale systems.

3.1.2 Multi-Level Operation

In addition to tailoring the system language to the class of modeling job, it must also be tailored to the type of user. In order to reach the largest population of potential users, a computerized system of decision aids would have to be developed to operate differently depending upon the skills and objectives of the user. We envision a hierarchy of decision aids having four basic levels of operation: Level I for the decision maker relatively unfamiliar with decision analysis and the operation of the computer system; level II for the "analyst" moderately experienced with the concepts of decision analysis; level III for the analyst with a thorough understanding of decision analysis methodology and good programming ability; and level IV for the "system programmer" who designs and extends the system's capabilities. Table 3.1 summarizes these four levels of system operation. The various levels of operation are hierarchically connected in the sense that one of the objectives of working in the higher levels (levels III and IV) is to expand the capabilities of the system at lower levels of operation (levels I and II). Since it will occasionally be desirable to build tailor-made computational aids to fit unique situations, the hierarchical structure of the system will provide the means for building new tools as needed. In this way, computerized tools can be developed for staff and line uses, strategic and tactical decisions, and planning and crisis environments.

Table 3.1
HIERARCHY OF AUTOMATED DECISION AIDS

LEVEL	PURPOSE	USER	REQUIRED USER EXPERIENCE	DECISION AID	USER-MACHINE INTERFACE
1	DECISION MAKING	EXECUTIVE	REQUIRES NO TRAINING IN DECISION ANALYSIS OR PROGRAMMING	<ul style="list-style-type: none"> -PROBLEM-ORIENTED ELICITATION AND ANALYSIS PROGRAMS -INTERACTIVE PROMPTING -PREPACKAGED MODELS FOR VARIOUS CLASSES OF DECISION PROBLEMS 	<ul style="list-style-type: none"> -INFORMATION ENTERED THROUGH INTERACTIVE GRAPHICS TERMINAL -GRAPHICAL DISPLAYS OF MODEL STRUCTURE, ELICITED INFORMATION, AND MODEL SOLUTIONS
2	ROUTINE DECISION ANALYSIS	ANALYST	REQUIRES GOOD UNDERSTANDING OF METHODOLOGY AND SOME PROGRAMMING EXPERIENCE	<ul style="list-style-type: none"> -PARAMETERIZED MODELS (DECISION TREES, NETWORKS, ETC.) -STRUCTURAL DEFINITION VIA OPERATOR RULES (HIERARCHIES) -ANALYST SELECTION OF MODELING FUNCTIONS 	<ul style="list-style-type: none"> -DATA INPUT THROUGH INTERACTIVE TERMINAL OR DATA FILES -ANALYST SELECTION OF OUTPUT FORMAT
3	ADVANCED DECISION ANALYSIS	ANALYST	REQUIRES THOROUGH UNDERSTANDING OF METHODOLOGY AND GOOD PROGRAMMING ABILITY	<ul style="list-style-type: none"> -DATA TYPE DEFINITION -DECISION TREES -NETWORK MODELS -COMPUTATIONAL GRAPHS -MODEL AND MODEL-CLASS DEFINITION -NON-PROCEDURAL PROGRAMMING -NON-DIRECTIONAL EQUATION PROCESSING -FUNCTIONAL LIBRARY 	<ul style="list-style-type: none"> -PROGRAMMING LANGUAGE FOR DATA TYPE DEFINITION
4	SYSTEM DESIGN AND IMPLEMENTATION OF MODELING LANGUAGE	SYSTEM-LEVEL PROGRAMMER	REQUIRES EXTENSIVE PROGRAMMING EXPERIENCE	<ul style="list-style-type: none"> -HOST LANGUAGE FOR DEFINING STYLE AND LINKING PROCEDURES -PROCEDURE DEFINITION LANGUAGE -SYSTEM EDITOR 	<ul style="list-style-type: none"> -SYSTEM COMPILER LANGUAGE

3.1.3 Analytical Capabilities for Decision Analysis

To be successful, of course, the language for any modeling system must be sufficiently general that the user can properly describe the behavior required of his model. It must contain an adequate range of data primitives, basic logical and arithmetic operations, conditional and unconditional branch facilities, an iterative capability, and provisions for the use of other software subsystems.

Comprehensive facilities for measuring the model's performance are an essential requirement. The system must contain a number of consistency checks to give the user confidence that the tools are being used properly. These checks will force the user to clarify his assumptions, verify the consistency of his preferences, and check his state of information. Certain basic calculations pertaining to decision analysis, such as sensitivity analysis, optimization functions, and probability processing techniques, are necessary. Since for larger problems some of these calculations may be beyond the computational capabilities of present generation computers, it may be useful to conceive of various levels of implementation for these computational abilities. Initially an effort would be made to implement only the lower levels of computational capabilities; the higher levels would be implemented as computer hardware and software technology progressed. Table 3.2 lists three levels of progressively more advanced computational capabilities for decision model design and analysis.

TABLE 3.2

CAPABILITIES OF MODELING AID

	LEVELS	1	2	3
ANALYTICAL FUNCTIONS				
<u>SENSITIVITIES</u>				
I/O SENSITIVITY (VALUE OF AN OUTPUT AS AN INPUT IS VARIED)		SENSITIVITY TO A SINGLE VARIABLE	JOINT SENSITIVITY TO TWO VARIABLES	MULTI-VARIABLE JOINT SENSITIVITY
TRADE-OFF SENSITIVITY (VALUE OF ONE INPUT AS ANOTHER INPUT IS VARIED HOLDING OUTPUTS FIXED)		SENSITIVITY OF ONE INPUT TO ANOTHER	SENSITIVITY OF ONE INPUT TO ANOTHER WITH CONSTRAINTS	
CHANGE RATES (PARTIAL DERIVATIVES)		DERIVATIVE OF OUTPUT TO INPUT, FIRST AND SECOND ORDER	DERIVATIVE OF OUTPUT TO SEVERAL INPUTS, TRADE-OFF DERIVATIVES	HIGHER ORDER DERIVATIVES
<u>OPTIMIZATION</u>				
OPTIMIZATION OF DETERMINISTIC MODEL		SINGLE DECISION VARIABLE	MULTIPLE DECISION VARIABLES	
OPTIMIZATION OF PROBABILISTIC MODEL		SINGLE DECISION VARIABLE	MULTIPLE DECISION VARIABLES, STATIC PROBLEMS	DYNAMIC DECISION PROBLEMS
<u>PROBABILITY PROCESSING</u>				
PROFIT LOTTERY GENERATION		SEVERAL INDEPENDENT PROBABILITY DISTRIBUTIONS	PROBABILISTIC DEPENDENCIES WITHOUT AUTOMATIC BAYES RULE PROCES- SING	ARBITRARY PROBABILISTIC DEPENDENCIES WITH AUTOMATIC BAYES RULE PROCESSING AND CONSISTENCY CHECKING
VALUE OF INFORMATION GIVEN FLEXIBILITY		PROBLEM STRUCTURE REDEFINED BY USER	VALUE OF INFOR- MATION FOR NON- DECISION-DEPENDENT VARIABLES	VALUE OF INFORMATION AND FLEXIBILITY WITH AUTOMATIC PROCESSING OF DECISION-DEPENDENT VARIABLES
PROXIMAL ANALYSIS		NO JOINT EFFECTS	CORRELATION AND INTERACTION AMONG VARIABLES	HIGHER ORDER APPROXIMATIONS

3.2 A GRAPHICAL LANGUAGE FOR MODEL DEVELOPMENT

One contribution of our research effort to develop a morphology for characterizing and analyzing decision problems has been the development of a graphical-algebraic language for decision model specification. The basic elements of this language, entities and operators, were defined in section 2.1. In that section we showed how entities and operators could be combined to represent system models. In this section we describe how this language may serve as the basis for a technique for model development. In addition, we illustrate how this technique might be implemented as a computerized system for decision analysis.

By far the most difficult and time consuming part of a decision analysis is the development of a decision model that captures the important elements of the problem. Even though the resulting model may be simple in form, considerable effort is likely to be required to choose problem variables and model parameters so as to represent adequately the essential aspects of the decision. Increasingly, though, resource allocation problems involve extremely complex systems. For example, the choice among pollution abatement technologies for an electric power plant requires the understanding of a system whose behavior patterns cannot be discerned from simplified approximations. Such systems may be quite large; their bounds are not determined by physical size, but rather by how much we need to know about the way the system interacts with its environment. For example, for the decision among sulfur oxide

pollution abatement strategies for an electric power plant, the system of interest is much larger than merely the power plant. Other elements which must be included in the system include the atmospheric conditions that govern the dispersal of the sulfur oxides and the resulting human population which becomes so exposed.

Thus, an important problem to be solved in the course of development of computer aids for decision making is to develop a language for compactly and efficiently deriving and expressing an appropriate model that describes the decision problem. The task of deriving a model of a system may be divided broadly into two subtasks: establishing the model structure and supplying the data. In Section 2.1 we described a method for concisely specifying model structure. The method was originally developed to represent the deterministic dependencies in continuous static models. However, we shall see in the latter part of this chapter that the method may be extended to apply to sequential systems containing discrete events as well. Few modeling techniques are sufficiently general that they can be applied to a broad cross-section of the model types discussed in Section 1.3.

Any system associated with a resource commitment problem will contain certain objects of interest. For the system describing sulfur oxide pollution from an electric power plant, objects of interest include the amount of electricity generated, the quantity of sulfur oxide emitted from the stacks, and the population downwind from the plant. Entities will be used to denote objects of interest

in a system. Model structure is defined by relationships which exist between the various system entities. An example would be the mathematical relationship that might exist between the quantity of electricity generated and power plant sulfur emissions. "Operators" will be used to represent such relationships among entities.

A useful property of the concepts of entities and operators was illustrated in Section 2.1; they allow us to express graphically model structure. The proposed technique of graphical representation has been applied to the system for analyzing sulfur oxide pollution abatement strategies from power plants.* About one-third of the resulting "branching structure" is illustrated in Figure 3.1.

In such branching structures, entities are represented graphically by blocks. The name or description of the entity is written within the block. Triangles represent operators and a line connecting an operator to an entity expresses the fact that the relationship among entities represented by the operator includes that entity. Entities nearer the top of the branching structure may be characterized as being more fundamentally important for judging the success or failure of a decision strategy. Entities nearer to the bottom of the structure are the more elemental properties of the system. In particular, those blocks which have no entering arrows are the most elemental system entities.

* This problem was the subject of a decision analysis recently completed by SRI for the National Academy of Science.

The lines between entities and operators shown in Figure 3.1 represent deterministic dependencies, and the arrows on these lines define a direction for computation. The rule satisfied by the connections between operators and entities is that given fixed values for all entities with arrows leading to an operator (input entities), a unique value is specified for all entities lying along arrows leading from that operator (output entities).

Elemental entities (those entities in Figure 3.1 with no entering arrows) may be classified as decision variables or as state variables. Decision variables are entities whose values represent decision alternatives and, hence, are set by the decision maker. In Figure 3.1 the elemental entities "Type coal consumed", "Coal washing technology", and "FGD (Flue Gas Desulfurization) technology" are decision variables representing possible means of sulfur oxide emission control available for the electric power plant. Decision variables are represented graphically by rectangular rather than circular blocks. The remaining elemental entities represent either parameters or state variables--entities whose values are determined by nature. Thus, the branching structure, when combined with the computational rule assignments for each operator, graphically represents the decision model's deterministic value structure. Given specific decision alternatives for each decision variable and specific values for all state variables, the value of the fundamental entity whose level determines the net worth of the decision outcome may be calculated.

An algorithm for constructing a branching structure such as that of Figure 3.1 will be given below, where we describe a computerized system designed to aid a decision maker in model development. Later, we shall show how an "influence diagram" may be combined with the branching structure for the deterministic value model. to represent dynamic learning effects in the model.

3.2.1 CADMUS (Computer Aids for Decision Makers: User System)

What follows is a tentative description of the operation of an interactive computer graphics system for decision making. We shall call this system CADMUS^{*}, an acronym for Computerized Aids for Decision Makers: User System. Although the characteristics which we ascribe to CADMUS are within the capabilities of present generation hardware and software technology, our objective is to communicate a perception of how such a system might usefully employ branching structures as a language for model building. Our objective is not to define the detailed characteristics of such a system. Before we become lost in the particulars of designing computer systems, we need to improve further our understanding of the alternative ways of describing decision problems.

To illustrate our perception of how CADMUS might use branching structures we shall use, as an example, the problem of choosing between sulfur pollution abatement strategies for fixed coal burning

* According to Greek mythology, Cadmus was a Phoenician prince reputed to have killed a dragon. Sowing the teeth of the dragon, Cadmus produced an army of armed men.

electric power plants.^{**} Due to the scope of the original analysis for this problem, a static model was chosen. State variables in this model are for the most part continuous. Branching structures are especially appropriate for representing the model structure for problems with these characteristics.

We assume that the system user in our example has some familiarity with decision analysis. This would put him somewhere in between levels 1 and 2 of user experience as outlined in Table 3.1.

A Brief Description of the Sulfur Problem

Sulfur oxide and particulate emissions have adverse consequences for human health and welfare, but the means for controlling these emissions entail considerable expense. The decision problem considered here concerns the choice between alternative sulfur emission abatement techniques for a representative 620MW coal burning electric power plant located in a rural area roughly 500km upwind from the New York Metropolitan area. Most of the sulfur content of coal is released upon combustion as sulfur dioxide, but atmospheric chemical processes may oxidize the sulfur dioxide to sulfuric acid aerosol and suspended particles of ammonium sulfate and other sulfate salts. Recent epidemiological data have indicated that these sulfates may give rise to serious and widespread health effects. Damage to material property from atmospheric sulfur oxides

^{**}As stated in a previous footnote, this problem was recently analyzed by the Decision Analysis Group at SRI. The results of this analysis were delivered to the U.S. Senate on March 3, 1975.

has been estimated to cause hundreds of millions of dollars in annual losses. "Acid rain" resulting from atmospheric sulfur oxides may lead to retarded growth in forests, deleterious effects on lakes and streams, damage to agricultural crops, and damage to building materials, statues and other material property. While SO_2 is invisible, sulfate particles do absorb and scatter light and, hence, may be partially responsible for visibility degradation.

The decision problem on emissions control is ultimately whether the owners of a power plant shall modify their operations by such means as installing a flue gas desulfurization (FGD) process, switching to a low sulfur content fuel, or installing a taller stack and intermittent control system. The adoption of the emission control strategy will result in higher costs to the owners of the power plant, and these higher costs will generally be passed on as higher prices to the consumers of electricity.

Therefore, obtaining an optimal emissions control strategy involves determining the cost-benefit trade-offs from using the various alternative means for control. The benefits from adopting the emissions control strategy come from the change in amount (and timing) of emissions into the atmosphere of sulfur oxide and other materials that may adversely affect human health, cause danger to other living organisms or material property, and result in effects such as visibility reduction that are aesthetically undesirable. The costs from adopting the strategy are the higher costs of

producing electricity. A decision between alternative strategies requires a balancing of the additional cost imposed on the generation of electric power against the value of emissions reduction. For the purposes of the analysis, the criteria for choosing among the alternative strategies was chosen to be minimization of total social cost to society.

3.2.2 Using CADMUS to Analyze the Sulfur Problem

Let us imagine that CADMUS, our computerized system of decision aids, has been built. How might it be used by a decision maker interested in modeling and analyzing the sulfur emissions problem? He begins by sitting down before his local computer driven terminal with video screen and throws a switch on the console indicating his request for CADMUS. A list of computer aids for decision analysis appears instantly on the screen, and our decision maker indicates his wish for assistance in model development by touching the light pen attached to his terminal to an appropriate item in this list.

Instantaneously, the video screen is transformed into the control panel illustrated in Figure 3.2. The center of the control panel, which displays messages from the computer system, may also be used as a "scratch pad" for drawing, typing, or writing. As the light pen is pulled across this area, a displayed "ink" track appears to flow from the pen. Items appearing here may be erased by scrubbing over them with the pen. Around the edges of the screen are various control "pushbuttons"; if one of these is "pushed" (by

EXPAND		<div> <div>DEPENDENT VARIABLE</div> <div>INDEPENDENT VARIABLE</div> <div>FROM</div> <div>TO</div> <div>STEP SIZE</div> </div>	<div> <div> <div>7</div> <div>4</div> <div>1</div> <div>0</div> </div> <div> <div>8</div> <div>5</div> <div>2</div> <div>.</div> </div> <div> <div>9</div> <div>6</div> <div>3</div> <div>*</div> </div> <div> <div>√</div> <div>1/x</div> <div>NOT</div> <div>#</div> </div> <div> <div>LOG</div> <div>SIN</div> <div>AND</div> <div>IF</div> </div> <div> <div>LN</div> <div>COS</div> <div>NO</div> <div>OR</div> </div> <div> <div>(</div> <div>)</div> <div>TAN</div> <div>YES</div> </div> <div> <div>x</div> <div>y</div> <div>LN</div> <div>YES</div> </div> </div>
CONTRACT			
ROLL UP			
ROLL DOWN			
ROLL LEFT			
ROLL RIGHT			
ERASE			
HARD COPY			
STORE PAGE			
RECALL PAGE			
DATA			
MODEL			
PLOT			
SENSITIVITY			
CHANGE RATE			
PROB DIST			
VAL OF INFO			

FIGURE 3.2 CONTROL PANEL FOR MODEL FORMULATION

touching the pen to it), the system performs the indicated action. The buttons in the lower right hand corner of the screen form a "calculator" for writing mathematical or logical equations. The buttons on the left of the screen are for system control.

The following question appears at the center of the video screen:

WHAT QUANTITY OR QUANTITIES WOULD YOU NEED TO KNOW TO EVALUATE THE OUTCOME OF YOUR DECISION?

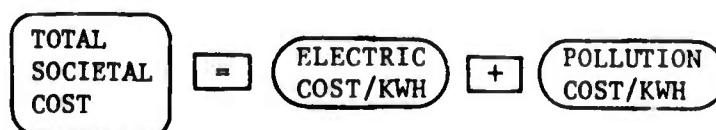
After thinking for a moment, our decision maker decides that he could evaluate a given emission abatement strategy if he could foretell the total social cost per kilowatt hour of electricity produced that the strategy would place on society. He, therefore, types TOTAL SOCIETAL COST/KWH on the computer terminal. CADMUS responds with another question:

SUPPOSE YOU HAD A CRYSTAL BALL. YOU MAY ASK IT ANY NUMERICAL QUESTION EXCEPT "WHAT IS 'TOTAL SOCIETAL COST/KWH'?" FOR WHAT QUANTITY OR QUANTITIES WOULD YOU ASK IN ORDER TO CALCULATE "TOTAL SOCIETAL COST/KWH"?

Our decision maker reasons that total societal costs will be the sum of the costs of electricity generation and pollution costs. He types ELECTRIC COST/KWH and the POLLUTION COST/KWH. Immediately the video screen represents these entities as entities in a graphical structure (Figure 3.3).

The relationship between the three entities is represented graphically by an "operator" dot, to which CADMUS has assigned the

number 1. To specify an analytic definition for the operator the decision maker addresses his attention to the calculator in the lower right hand corner of the screen. Thinking "total societal cost will be the sum of electric cost and pollution cost," he uses his light pen to write *



by alternately touching the appropriate entity blocks in the branching structure and operator buttons on the calculator. As he proceeds, the equation defining operator number one appears at the bottom of the video screen.

The decision maker touches the light pen to the entity block marked "pollution cost/kwh" and CADMUS answers:

SUPPOSE YOU HAD A CRYSTAL BALL. YOU MAY ASK ANY NUMERICAL QUESTION EXCEPT "WHAT IS 'POLLUTION COST/KWH'?" FOR WHAT QUANTITY OR QUANTITIES WOULD YOU ASK IN ORDER TO CALCULATE "POLLUTION COST/KWH"?

"Well," thinks our decision maker, "I can calculate the number of pounds of sulfur emitted from the stacks per kilowatt hour, so if I knew the pollution cost per pound of sulfur emitted I could cal-

As mentioned above, a user may "push" a "button" illustrated on the video screen by touching that button with the light pen. In the text we represent the pushing of a button by enclosing the button label within a rectangular figure. For permanent control buttons the rectangular figures have right-angled corners. User defined entity buttons are enclosed within figures with rounded corners. Rectangular blocks are not placed around those items that are entered through the terminal keyboard.

culate the pollution cost per kilowatt hour." He types POLLUTION COST/LB SULFUR EMITTED and LBS SULFUR EMITTED/KWH, causing the branching structure to disappear momentarily then reappear as shown in Figure 3.4. Using the light pen, the entity blocks in the branching structure, and the "calculator" in the left hand corner of the screen, our decision maker "pushes" the right buttons so as to express pollution cost per kilowatt hour as the product of pollution cost per pound of sulfur emitted and the number of pounds of sulfur emitted per kilowatt hour. The equation appears near the bottom of the screen as shown in Figure 3.4.

Next our decision maker touches his light pen to the block marked "electric cost/kwh". The familiar question, this time asking him what he would need to know to define "electric cost/kwh," appears on the screen. The decision maker types CAPITAL COSTS, OPERATING COSTS, and FUEL COSTS and uses the calculator to define "electric cost/kwh" as the sum of these quantities. The video screen now appears as in Figure 3.5.

At this point it occurs to our decision maker that he may not have defined his system broadly enough. Since sulfur oxides and other pollutant emissions from power plants are a by-product of electricity generation, one alternative for reducing these emissions is to reduce the electric output (the load factor) for the plant. Therefore, the number of kilowatt hours generated by the plant may be a variable. "The ultimate determinant of the success of a given strategy will be the total annual costs imposed on society," thinks

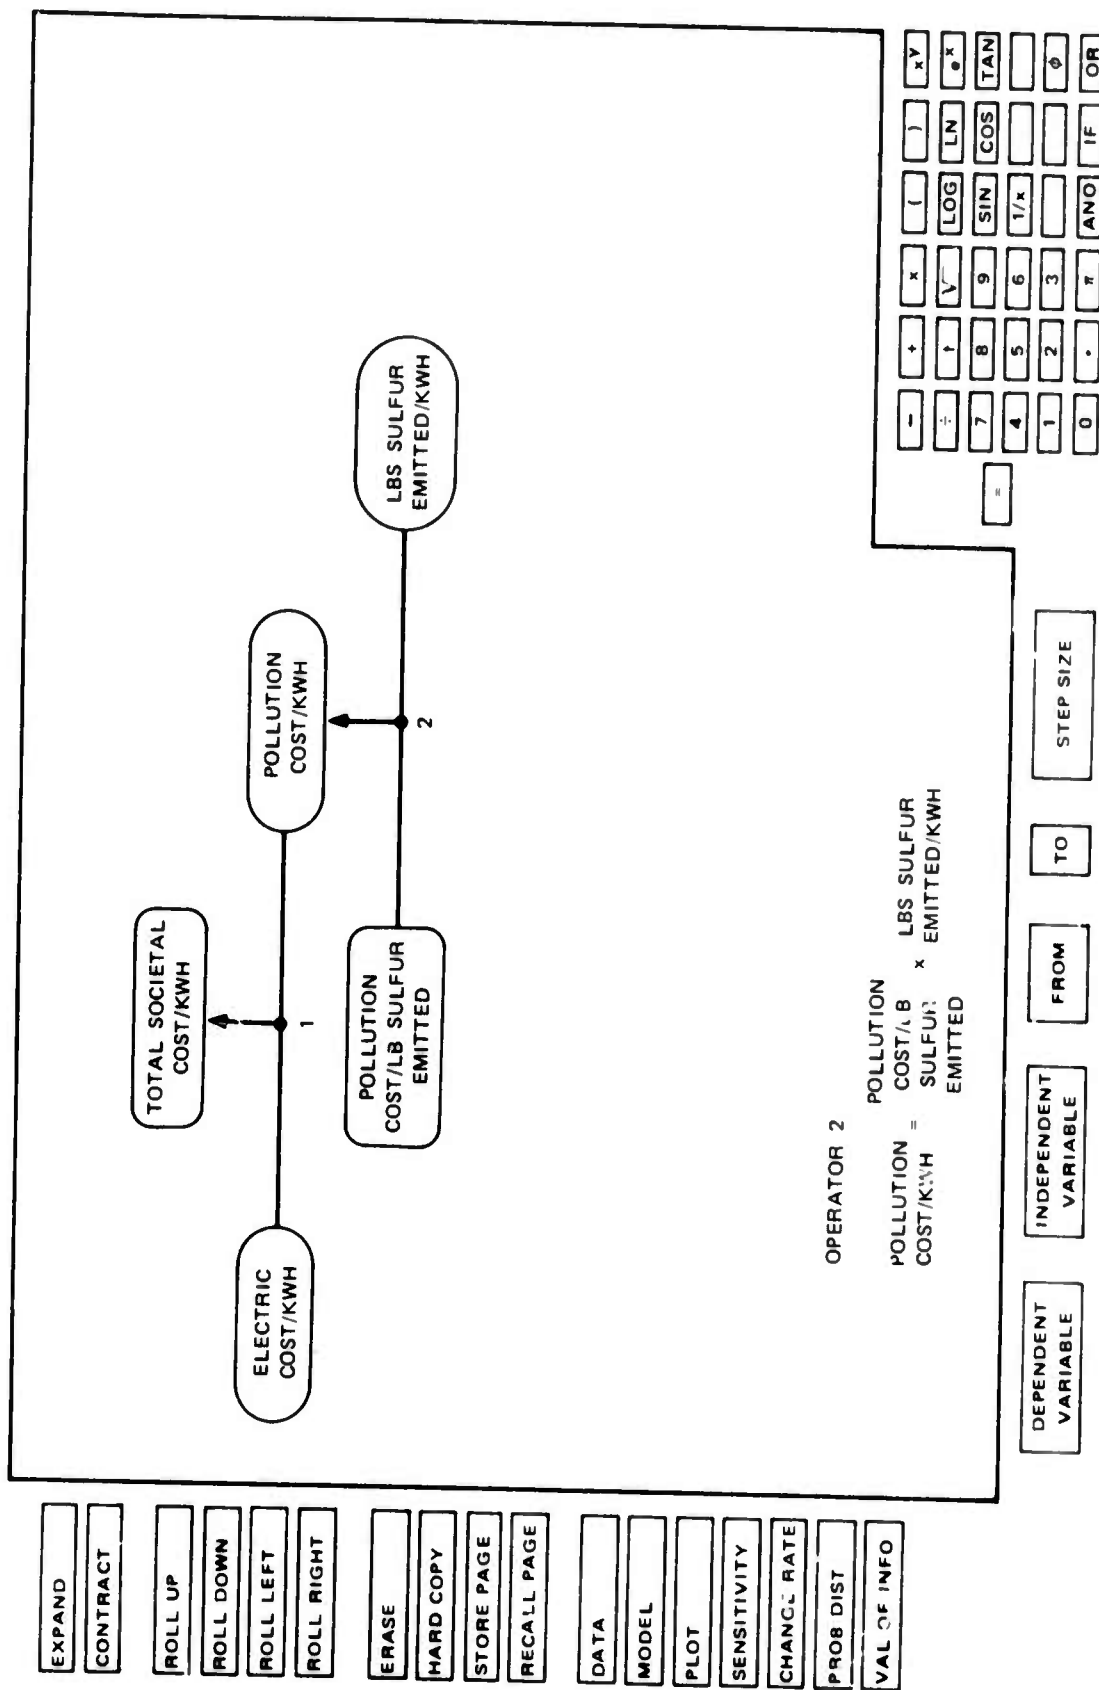


FIGURE 3.4 STEP TWO FOR MODELING THE SULFUR EMISSION PROBLEM

our decision maker, and he types

TOTAL ANNUAL SOCIETAL COST $\boxed{=}$ ANNUAL KWH GENERATED X

TOTAL
SOCIETAL
COST/KWH

CADMUS labels this equation as the definition of operator number 4, recognizes that two new entities have been defined, and locates the proper place in the branching structure to add the additional entities, operator and connecting lines. An instant later the branching structure appears as shown in Figure 3.6.

Our decision maker may continue to use the branching technique until he has defined all entities in terms of terminal, elemental quantities that he feels comfortable estimating. Operator definitions need not be confined to simple arithmetic functions like addition and multiplication. If necessary more general functions or subroutines may be used to define the relationships among system entities. For example, in the actual branching structure which was used for the sulfur emissions problem (partially illustrated in Figure 3.1), the operator relating the entity block "increase in SO₄ concentration" to the entities below it was defined by a differential equation.

Representation of the Deterministic Value Model

For the purposes of our example, let us suppose that our decision maker continues expanding the system by branching from entities until the structure appears as shown in Figure 3.7. At this point our decision maker feels that with a little research he can

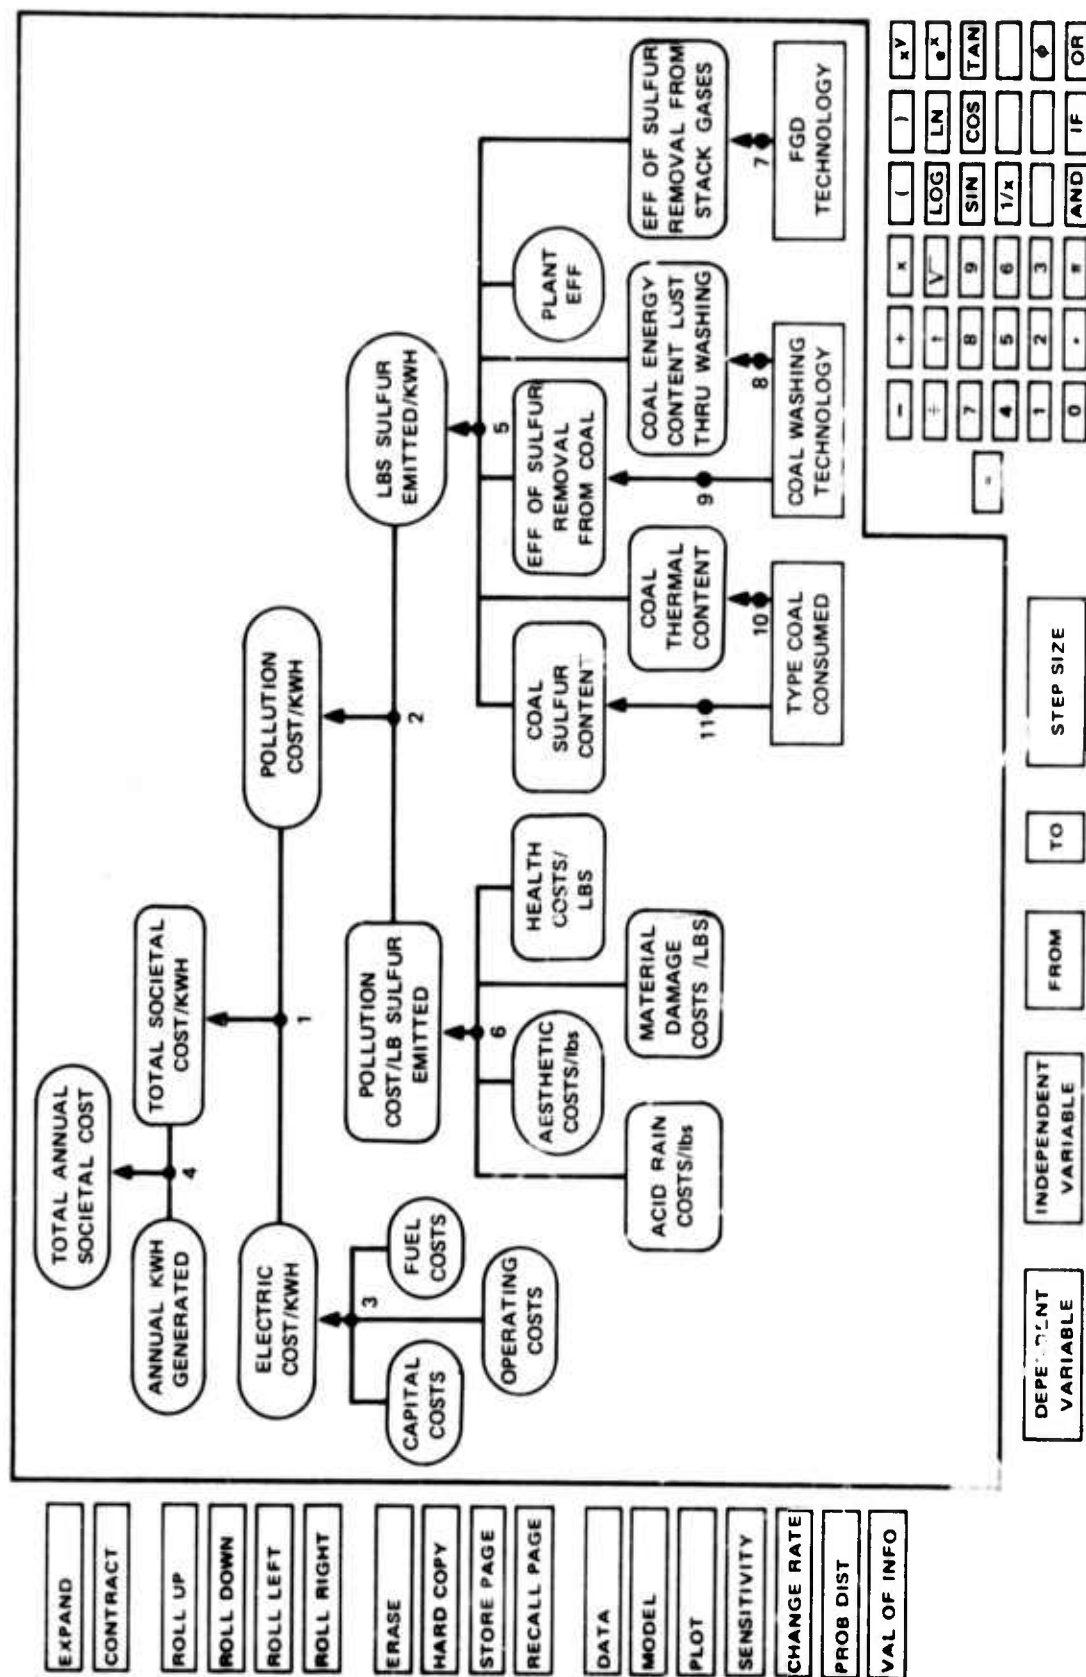


FIGURE 3.7 STEP TEN FOR MODELING THE SULFUR POLLUTION PROBLEM

come up with estimates of values for the elemental state variable entities, and he sees that the important decision alternatives "type coal consumed," "coal washing technology," and "FGD (flue gas desulfurization) technology" have all been represented.

The branching structure shown in Figure 3.7 (and for that matter each structure that appeared in the development of Figure 3.7) when combined with analytic operator definitions completely specifies the deterministic structure necessary for evaluating various decision strategies. Given specific alternatives for the decision variables and given specific values for the elemental state variables, evaluation of the branching structure will produce a value for total annual societal cost. The branching structure, therefore, represents a mathematical model for the sulfur emission problem. An important question is, "How good is this model?" CADMUS supplies a number of tests for model evaluation.

Problem Analysis

One important capability possessed by CADMUS is the ability to perform sensitivity analyses. In sensitivity analysis, the decision analyst tries to determine the change in the model's selection of alternative actions or outcome values that would result from a given change in the model's assumptions. Assumptions that produce small changes are apparently relatively insignificant, while assumptions that produce considerable changes are likely quite significant. CADMUS provides for a number of automatic sensitivity calculations. For example, suppose our decision

maker felt uneasy about his estimates of the various pollution costs per pound of sulfur emitted. He might feel that these costs could add up to anything between zero and as much as eighty cents per pound of sulfur emitted. Knowing how much this uncertainty is reflected as uncertainty in, for example, total societal cost/kwh will be important information in determining whether further effort should be expended in clarifying the estimation of pollution costs.

Suppose, therefore, that our decision maker elects to calculate the way the value of total societal cost/kwh behaves as pollution cost/lb of sulfur emitted is varied. He begins by choosing specific alternatives for the decision variables:

TYPE COAL CONSUMED = HIGH SULFUR COAL

COAL WASHING TECHNOLOGY = NONE

FGD TECHNOLOGY = NONE

Control buttons for the various alternatives will have been defined when the operators 7, 8, 9, 10, and 11 in Figure 3.7 are defined. (This process will be illustrated in the following section where we apply the branching structure technique to a discrete event problem.) Next nominal values are assigned to elemental state variables:

PLANT EFFICIENCY = 33 %

CAPITAL COSTS = \$ 250 /KW

OPERATING COSTS = 1 MILLS/KWH

FUEL COSTS = 10.3 MILLS/KWH. *,**

By touching his light pen to the SENSITIVITY control button, the decision maker signals CADMUS that he wishes to perform a sensitivity calculation. Pushing DEPENDENT VARIABLE and then TOTAL SOCIETAL COST/KWH establishes the dependent variable in the relationship. Our decision maker wishes to learn how the value of this variable varies as he varies pollution cost/lb sulfur emitted between the values of zero and 80¢. In order to communicate this to CADMUS he uses his light pen to push the respective buttons INDEPENDENT VARIABLE and POLLUTION COST/LB SULFUR EMITTED and then FROM 0 TO .80. The video screen momentarily goes blank and then appears as Figure 3.8, showing the functional

*To simplify notation, numbers which require a sequence of button pushing operations will be represented as if they could be specified by a single button. For example, the notation 10.3 means that the buttons 1 0 . 3 are pushed in succession.

**Actually these values will be functions of the decision variables. Indeed, such functional connections were made in the branching structure that was actually used in the analysis for this problem. Such connections have not been shown in Figure 3.1 so as to simplify the discussion.

- EXPAND
- CONTRACT
- ROLL UP
- ROLL DOWN
- ROLL LEFT
- ROLL RIGHT
- LRASE
- HARD COPY
- STORE PAGE
- RECALL PAGE
- DATA
- MCDEL
- PLOT
- SENSITIVITY
- CHANGE RATE
- PROB DIST
- VAL OF INFO

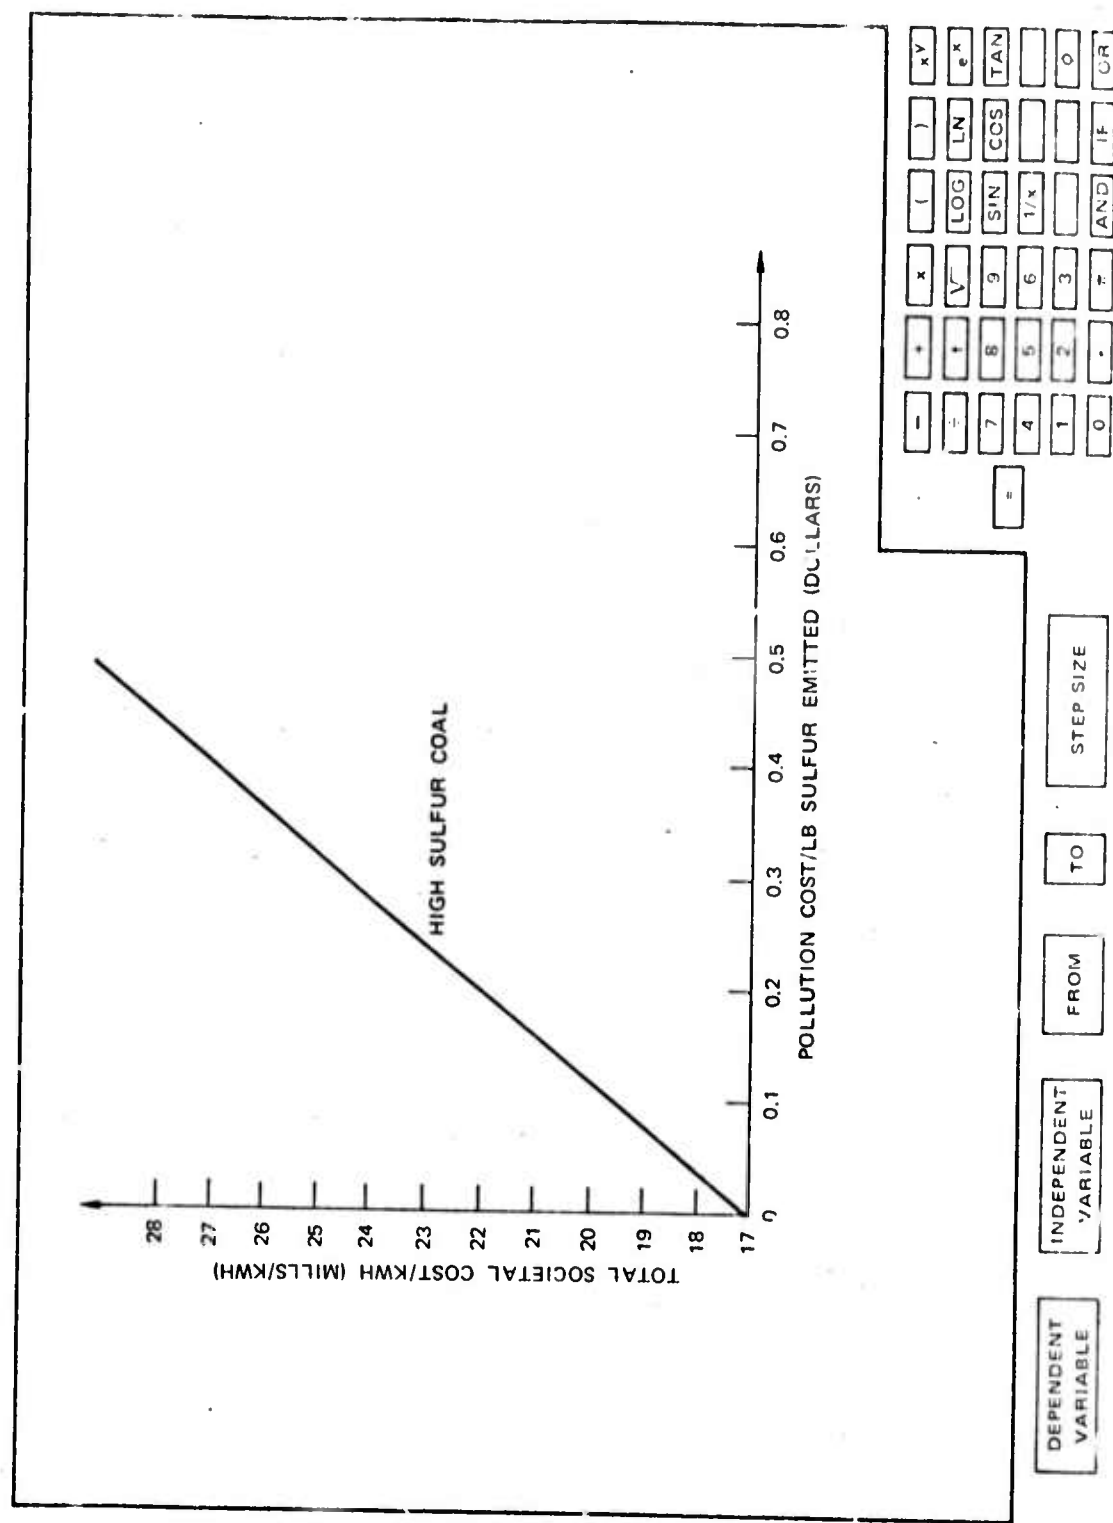


FIGURE 3.8 SENSITIVITY OF TOTAL SOCIETAL COST/KWH TO POLLUTION COST/LB SULFUR EMITTED (HIGH SULFUR COAL)

relationship established between the variables "total societal cost/kwh" and "pollution cost/lb sulfur emitted."* Since "total societal cost/kwh" is the sum of "electric cost/kwh" and the product of "lbs sulfur emitted/kwh" and "pollution cost/lb sulfur emitted," the relationship is indeed linear as shown.

Our decision maker may also wish to see similar sensitivity plots for other decision strategies. For example, if he were to set "type coal consumed" to "low sulfur eastern coal" and re-ran the sensitivity, the result would appear something like that shown in Figure 3.9. Similarly, if he reset "type coal consumed" to "high sulfur coal" and set "FGD technology" to "FGD," the result of the sensitivity run would appear something like that shown in Figure 3.10. An instructive exercise would be to display simultaneously the various sensitivity plots for the various decision alternatives. By storing the results and then pushing RECALL, our decision maker could generate the plot shown in Figure 3.11. Such a plot illustrates the closed loop sensitivity of "total societal cost/kwh" to "pollution cost/lb sulfur emitted." The closed loop sensitivity shows how the optimal value of the dependent variable changes with different values for the independent variable. Since the objective is to choose the decision strategy so as to minimize "total societal cost/kwh," this sensitivity is given by the piecewise linear curve formed by the lower envelope of the three straight line curves.

*The numerical values shown in Figure 3.8 and those that follow are meant to be illustrative only.

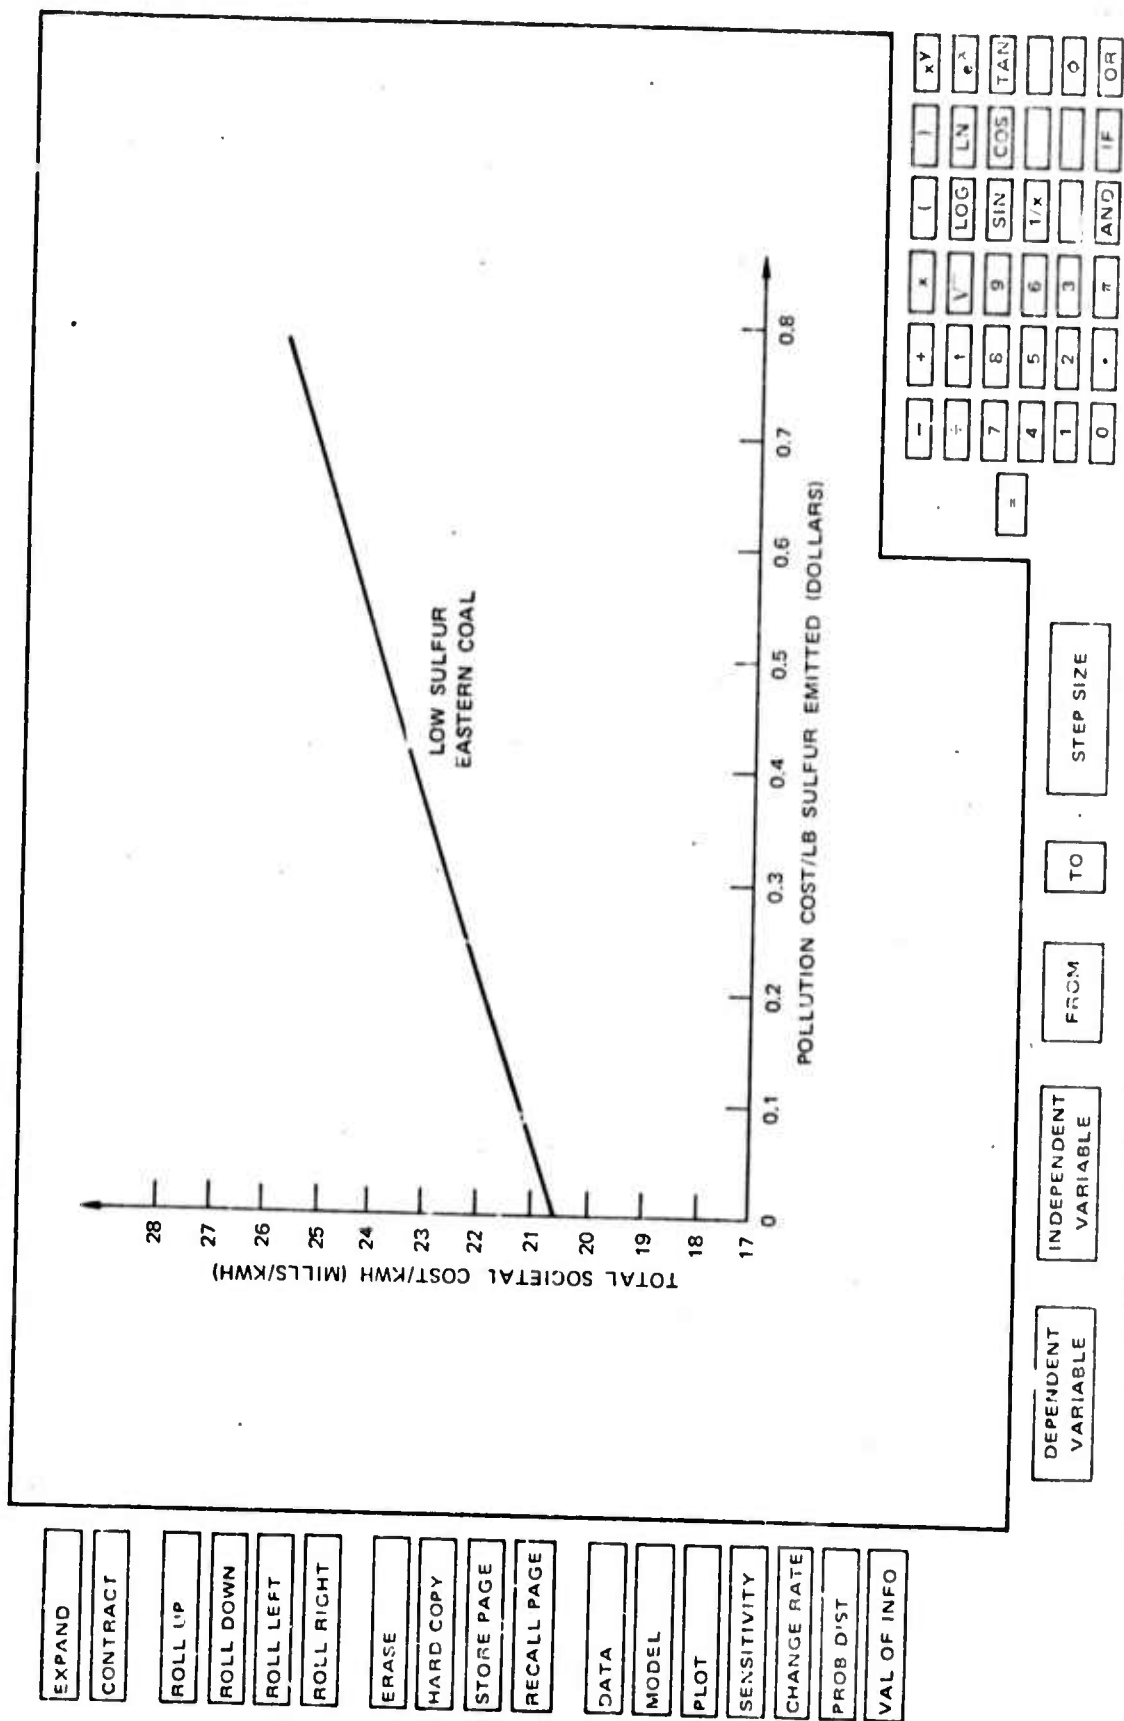


FIGURE 3.9 SENSITIVITY OF TOTAL SOCIETAL COST/kWh TO POLLUTION COST/LB SULFUR EMITTED (LOW SULFUR EASTERN COAL)

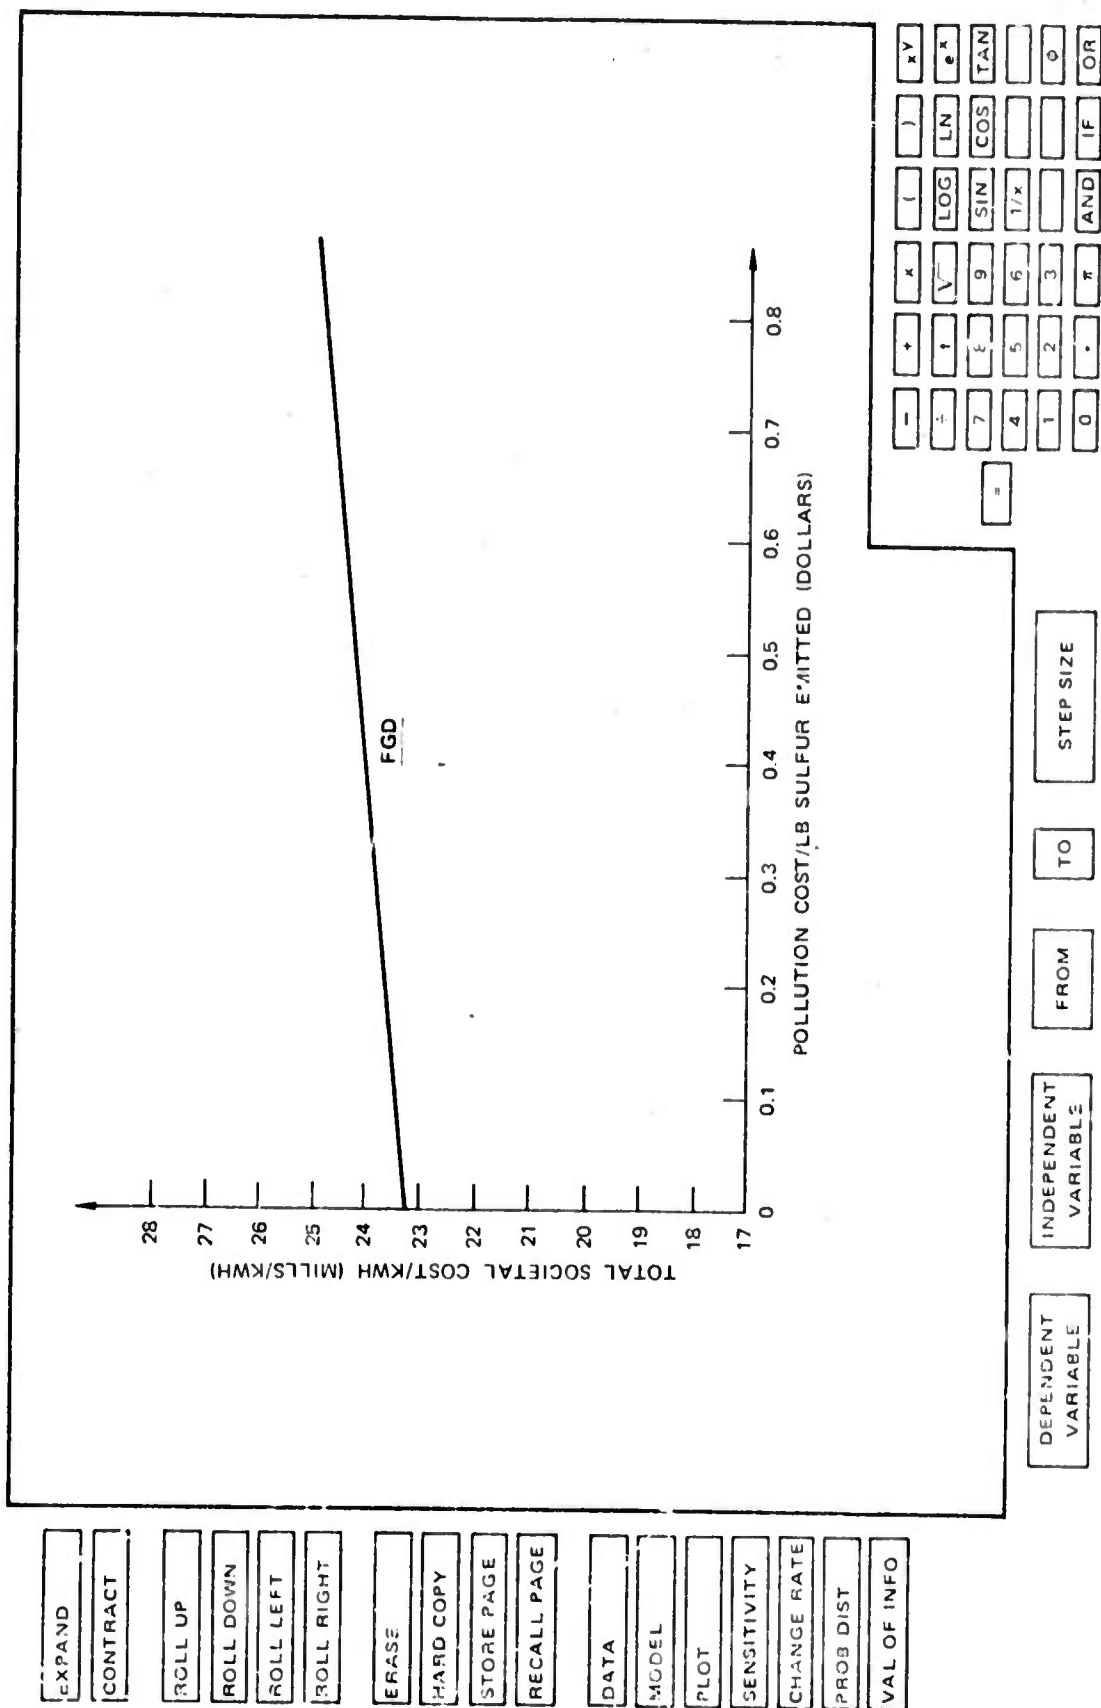


FIGURE 3.10 SENSITIVITY OF TOTAL SOCIETAL COST/KWH TO POLLUTION COST/LB SULFUR EMITTED (FLUE GAS DESULFURIZATION)

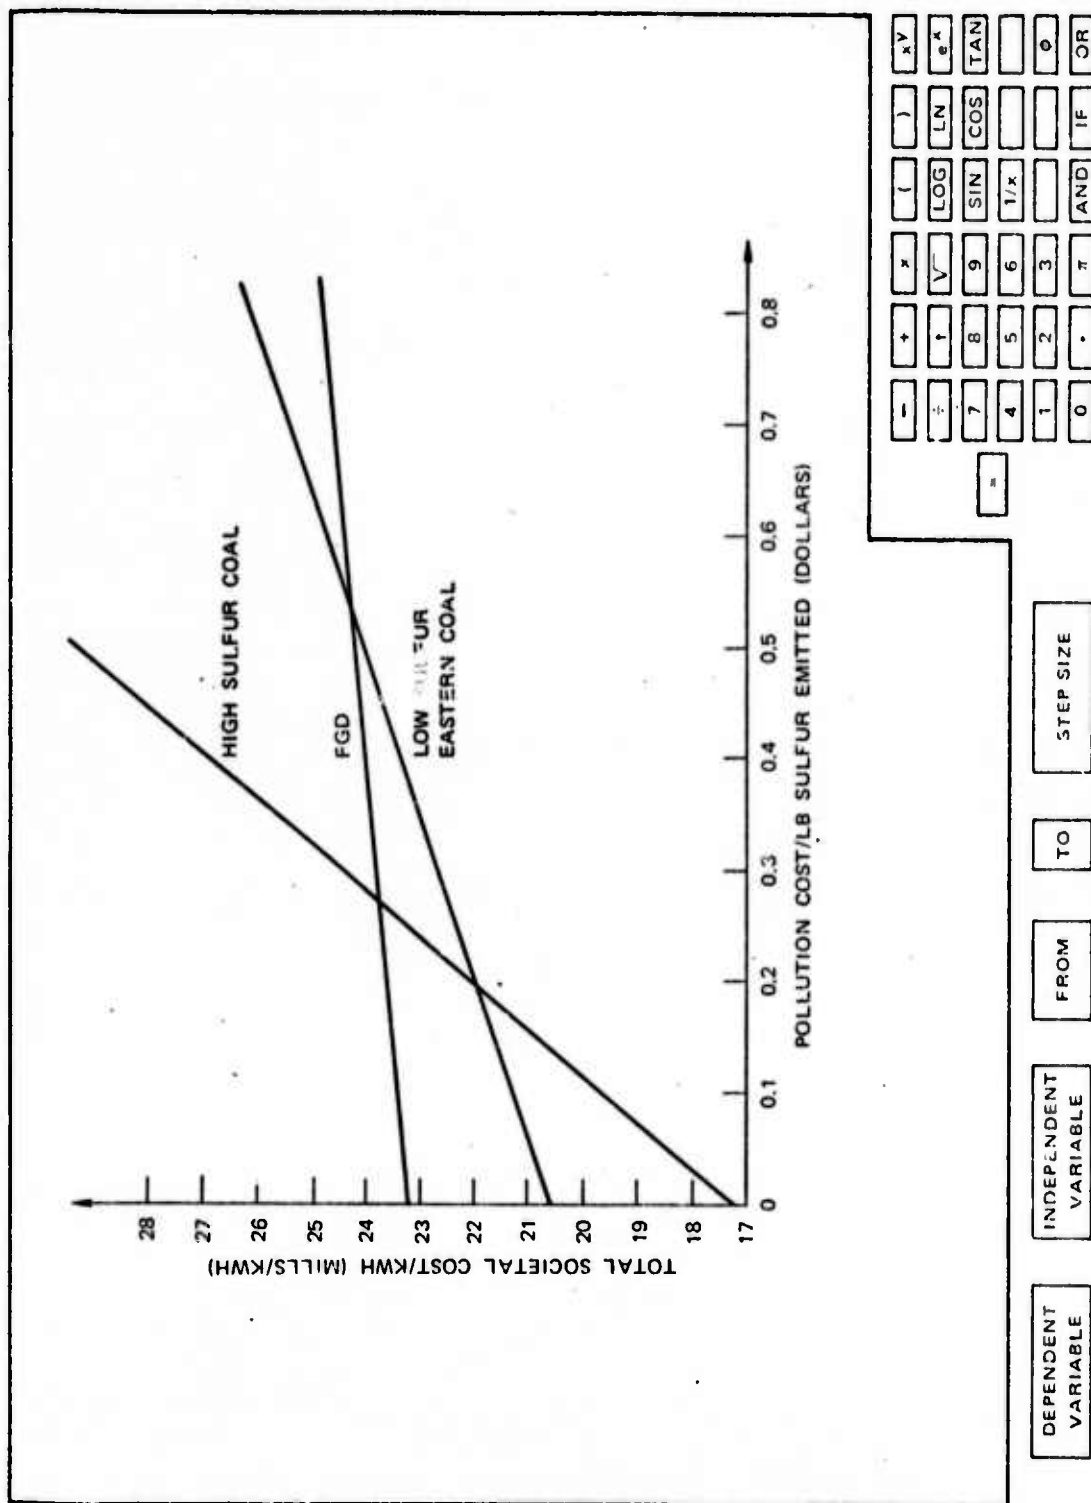


FIGURE 3.11 SENSITIVITY OF TOTAL SOCIETAL COST/KWH TO POLLUTION COST/LB SULFUR EMITTED

CADMUS may be requested to perform a number of additional calculations which will enable the decision maker to evaluate his alternative decision strategies. For example, he may ask to check for dominated courses of action; that is, decision strategies which for all possible state outcomes yield values equaled or exceeded by other alternative strategies. The decision maker may request to see the "profit lottery" associated with a given course of action. CADMUS will then supply the cumulative probability distribution which will give the probability that the outcome associated with that decision strategy will have a value less than or equal to any given amount.

To illustrate the latter possibility suppose that our decision maker were to specify a probability distribution for health cost/lb of sulfur emitted such as that shown in Figure 3.12. A number of methods to aid the decision maker in assessing such a distribution are provided by CADMUS. By touching his light pen to **PROB DIST** ("probability distribution" control button) and then to **TOTAL ANNUAL SOCIETAL COST**, the decision maker would instruct CADMUS to produce cumulative probability distributions for "total annual societal cost" under the various decision alternatives (Figure 3.13). The curve associated with a given decision strategy would show, for that strategy, the probability that total annual societal cost lay below any given amount.

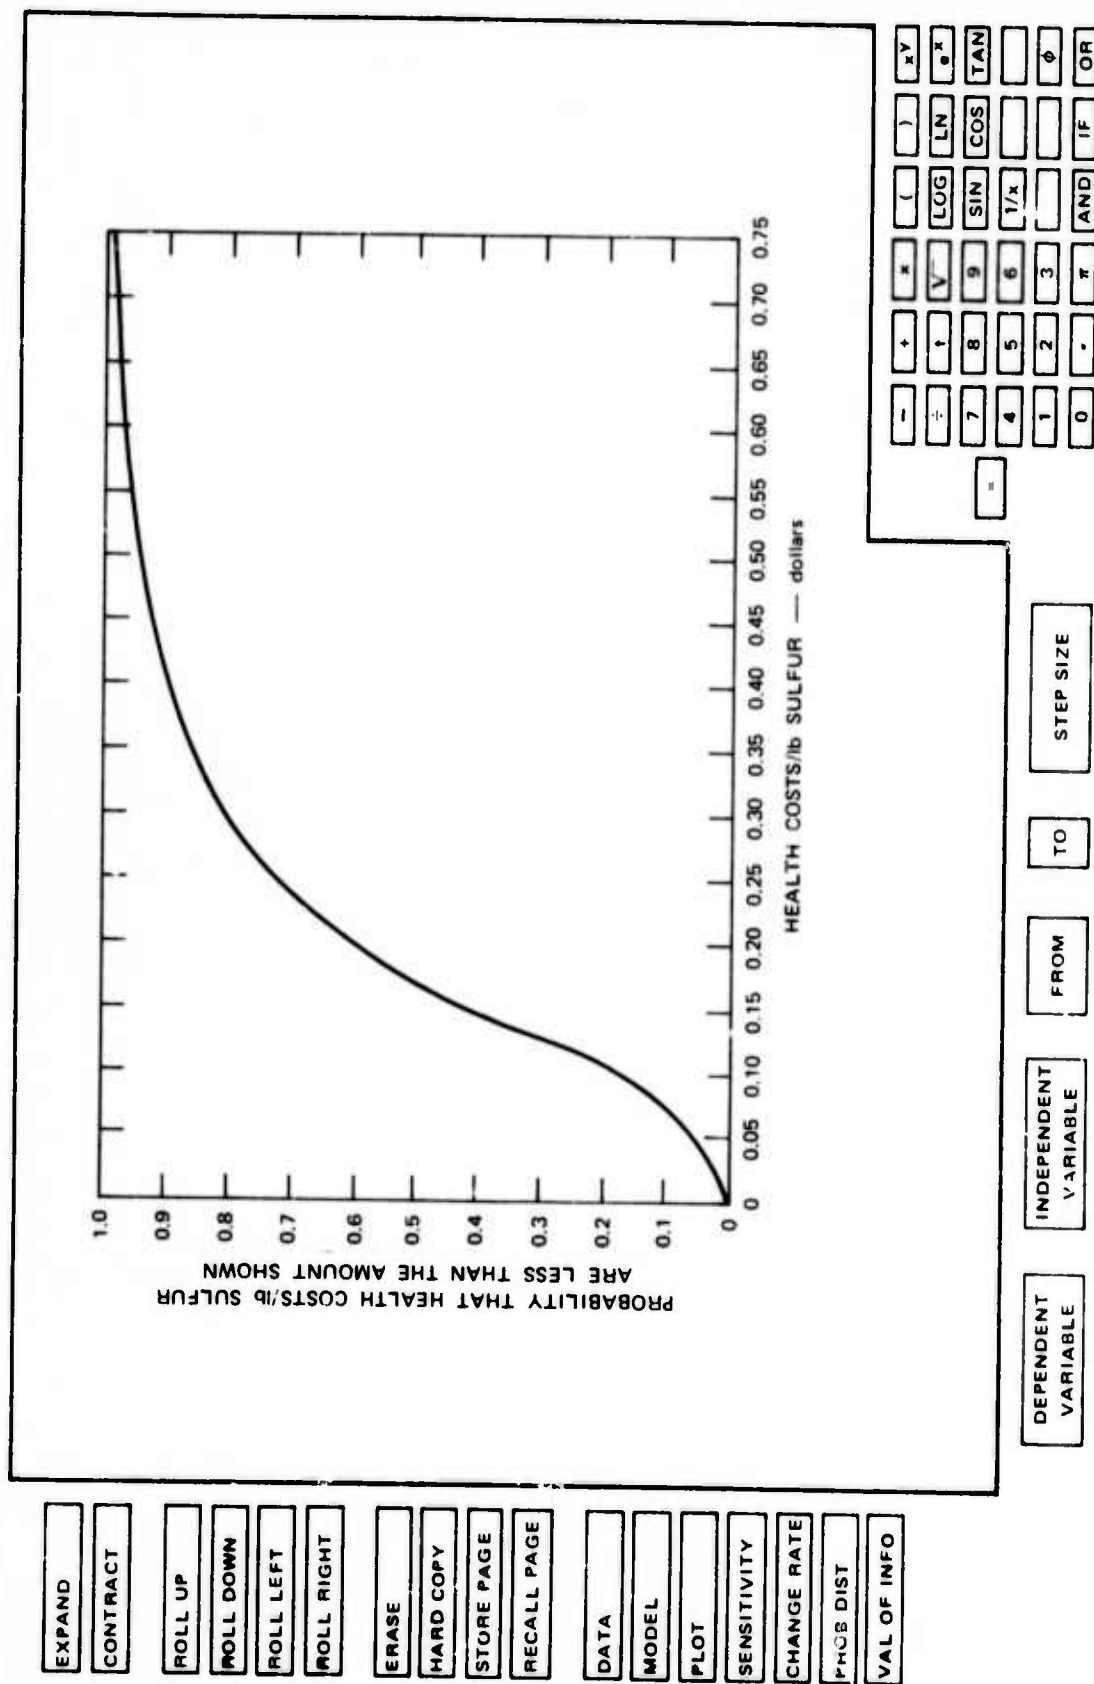


FIGURE 3.12 PROBABILITY DISTRIBUTION ILLUSTRATING UNCERTAINTY IN HEALTH COSTS PER POUND OF SULFUR EMISSIONS

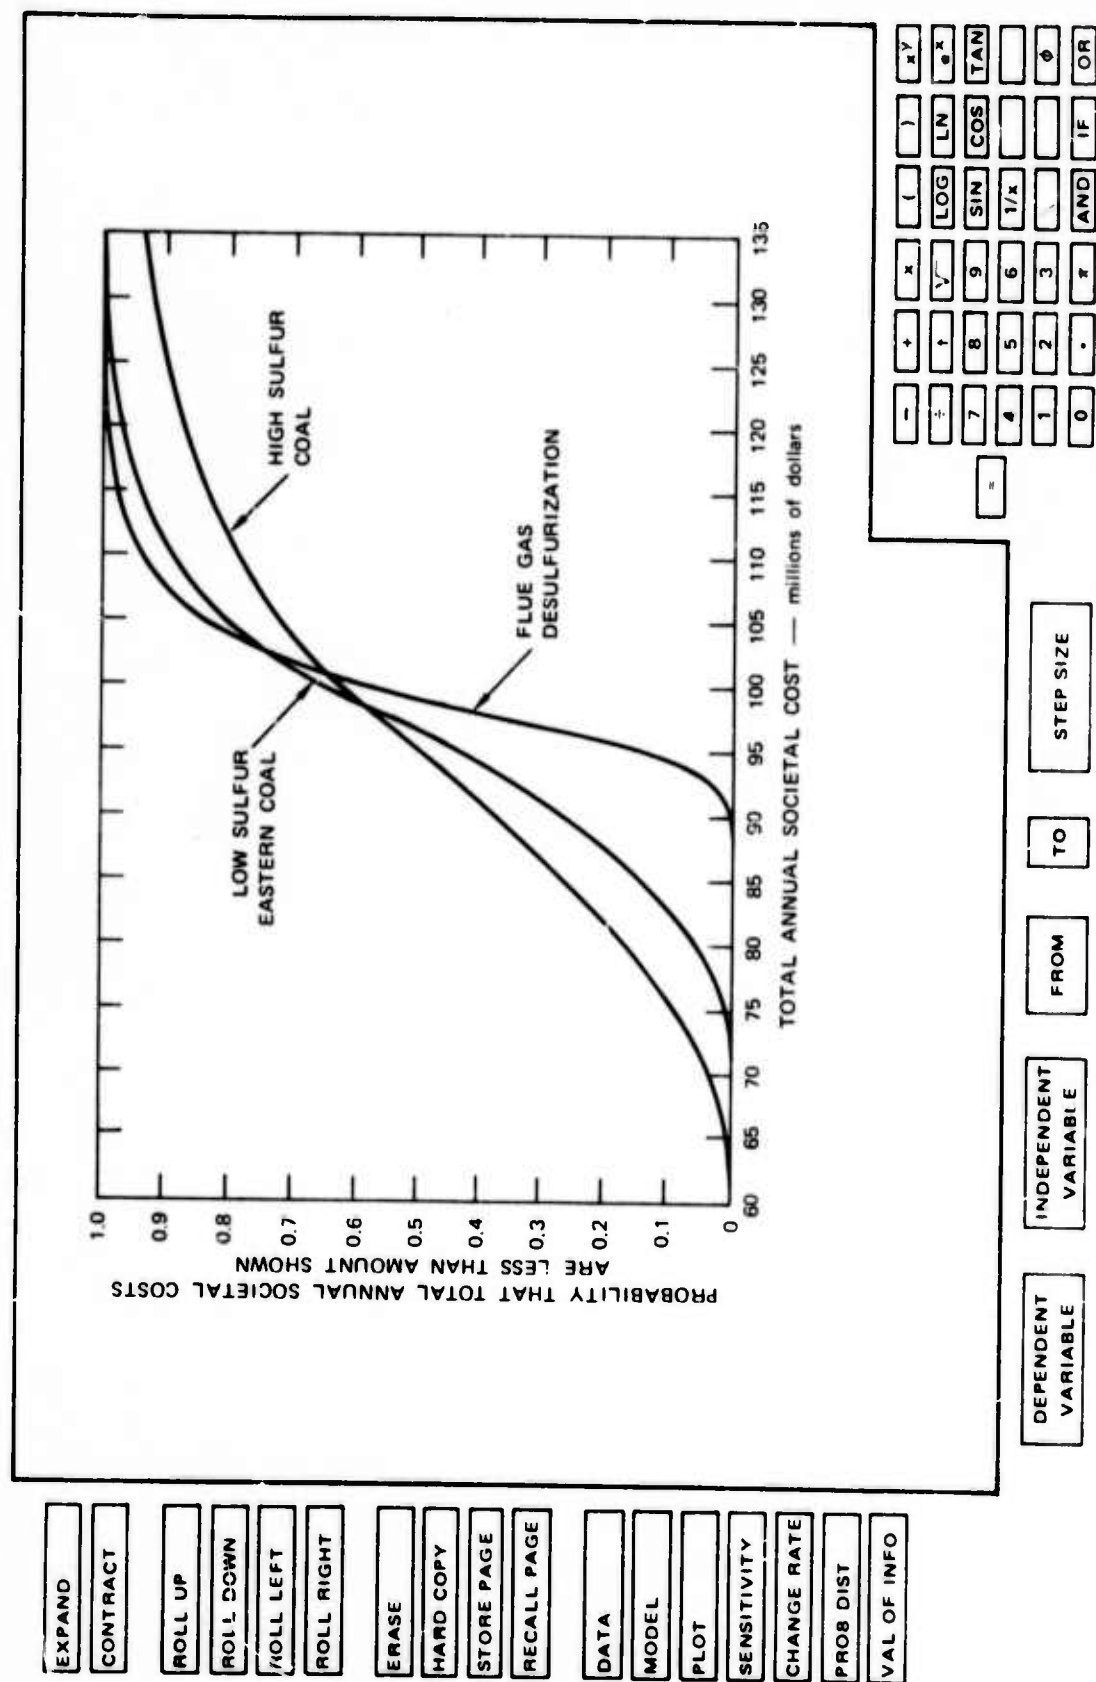


FIGURE 3.13 PROBABILITY DISTRIBUTION ILLUSTRATING UNCERTAINTY IN TOTAL ANNUAL SOCIETAL COSTS FROM 620 MW COAL BURNING ELECTRIC POWER PLANT

3.2.3 Conclusions Concerning the Applicability of the Branching Structure Approach to Single Stage Decision Models.

In our preceding exercise the postulated computerized system for decision analysis was employed primarily to illustrate the usefulness of the branching structure language of entities and operators for representing single stage deterministic models. The branching structure language provides a convenient method to record a well-specified deterministic value model. The language is also a tool to aid in model construction. Using the "crystal ball" algorithm a user can develop a model right at the computer terminal.

3.3 APPLICATION OF THE GRAPHICAL LANGUAGE FOR MODEL DEVELOPMENT TO A SEQUENTIAL PROBABILISTIC DECISION PROBLEM WITH A DISCRETE EVENT STRUCTURE.

The concept of using a branching structure as a modeling language was originally developed to apply to deterministic, single stage continuous variable models. At the end of the previous section we showed that relaxation of the deterministic restriction is not difficult. Basically, by holding the decision variables at a specific setting and specifying probability distributions on the state variables it is straightforward to generate a "profit lottery," a probability distribution on the output. Surprisingly, it is also possible to relax the other two constraints and consider sequential decisions and discrete variables within the same general branching structure.

To illustrate, we shall use a sample problem called the Used Car Buyer, Howard [4a]. The Used Car Buyer, a decision problem of

moderate difficulty, is used at Stanford University to teach decision analysis to graduate students.

Our decision maker, named Joe, is interested in purchasing a used car. He has found a car he likes, but because of uncertainties involving possible repair costs, he is unsure about whether he should buy or refuse the deal offered by the salesman. The car might be a "peach" or it might be a "lemon". If it is a peach it will have a defect that will require Joe to repair exactly one of the car's ten major systems. If it is a lemon Joe will soon have to repair exactly six of the car's major systems. The salesman offers to let Joe take the car to his mechanic to check out the condition of one or more of the car's ten systems. In addition, he offers Joe the possibility of purchasing a guarantee plan (an "anti-lemon feature"), which will help pay for any repair costs that might arise after purchase.

Suppose that Joe decides to seek help for his problem from CADMUS. He sits down before his terminal, throws a switch, and asks CADMUS for help in model formulation. CADMUS asks:

WHAT QUANTITY OR QUANTITIES WOULD YOU NEED TO KNOW TO
EVALUATE THE OUTCOME OF YOUR DECISION?

Accustomed to thinking in monetary terms, Joe types NET PROFIT into the computer. CADMUS responds with:

SUPPOSE YOU HAD A CRYSTAL BALL. YOU MAY ASK IT ANY NUMERICAL
QUESTION EXCEPT "WHAT IS 'NET PROFIT'?" FOR WHAT QUANTITY
OR QUANTITIES WOULD YOU ASK IN ORDER TO CALCULATE "NET PROFIT"?

Joe enters VALUE and then COST. Immediately the video screen represents these entries as entities in a graphical structure (Figure 3.14). Joe now addresses his attention to the calculator in the lower right hand corner of the screen. Thinking "net profit equals the difference between revenue and costs", Joe uses his light pen to "punch" in

$$\text{NET PROFIT} = \text{VALUE} - \text{COST}$$

by alternately touching the appropriate entities in the branching structure and operator buttons on the calculator. As he proceeds, the equation defining operator number one appears at the bottom of the video screen.

Joe touches the light pen to the entity block marked "cost" and CADMUS answers:

SUPPOSE YOU HAD A CRYSTAL BALL. YOU MAY ASK ANY NUMERICAL QUESTION EXCEPT 'WHAT IS 'COST'?' FOR WHAT QUANTITY OR QUANTITIES WOULD YOU ASK IN ORDER TO CALCULATE "COST"?

"Well," thinks Joe as he types on the terminal, "my costs will consist of the asking PRICE, a REPAIR COST if the car breaks down, a GUARANTEE COST if I take advantage of the guarantee plan, and a TEST COST if I ask my mechanic to look over the car before I buy it." The branching structure on the video screen now appears as in Figure 3.15. Using the calculator, Joe has no difficulty "pushing" the right buttons to express cost as the sum of price, repair cost, guarantee cost, and test cost.

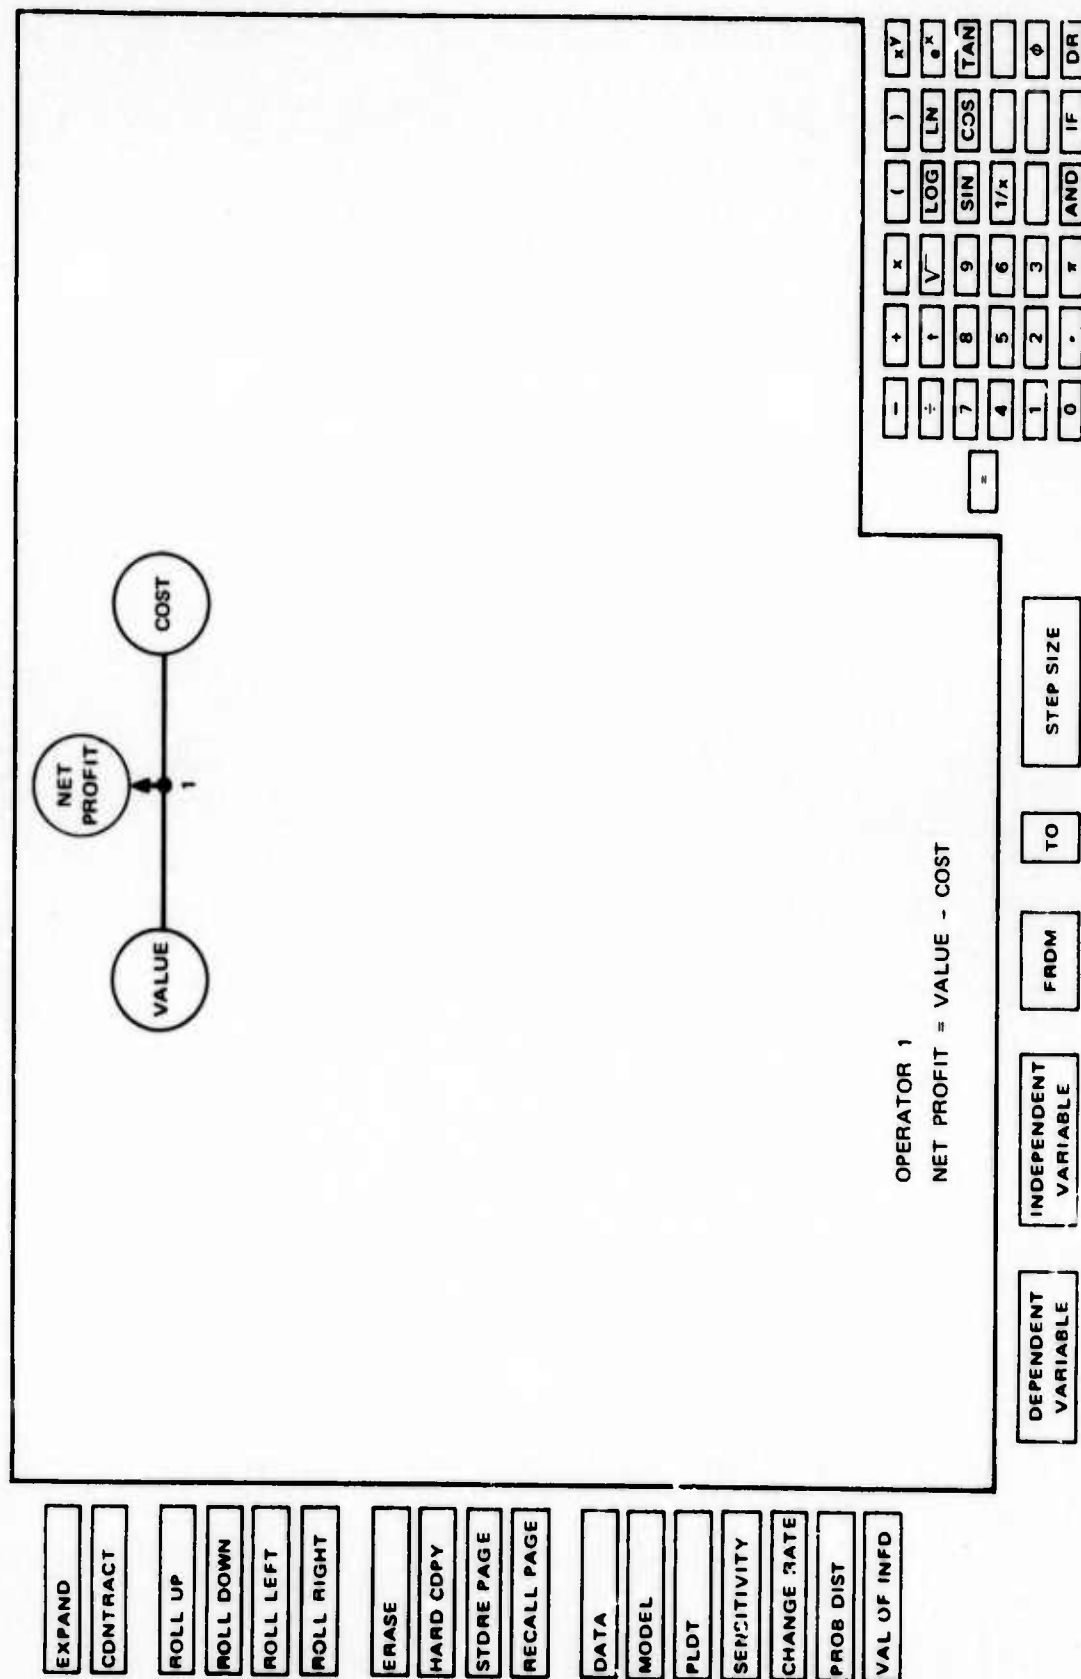


FIGURE 3.14 STEP ONE FOR MODELING THE USED CAR BUYER PROBLEM

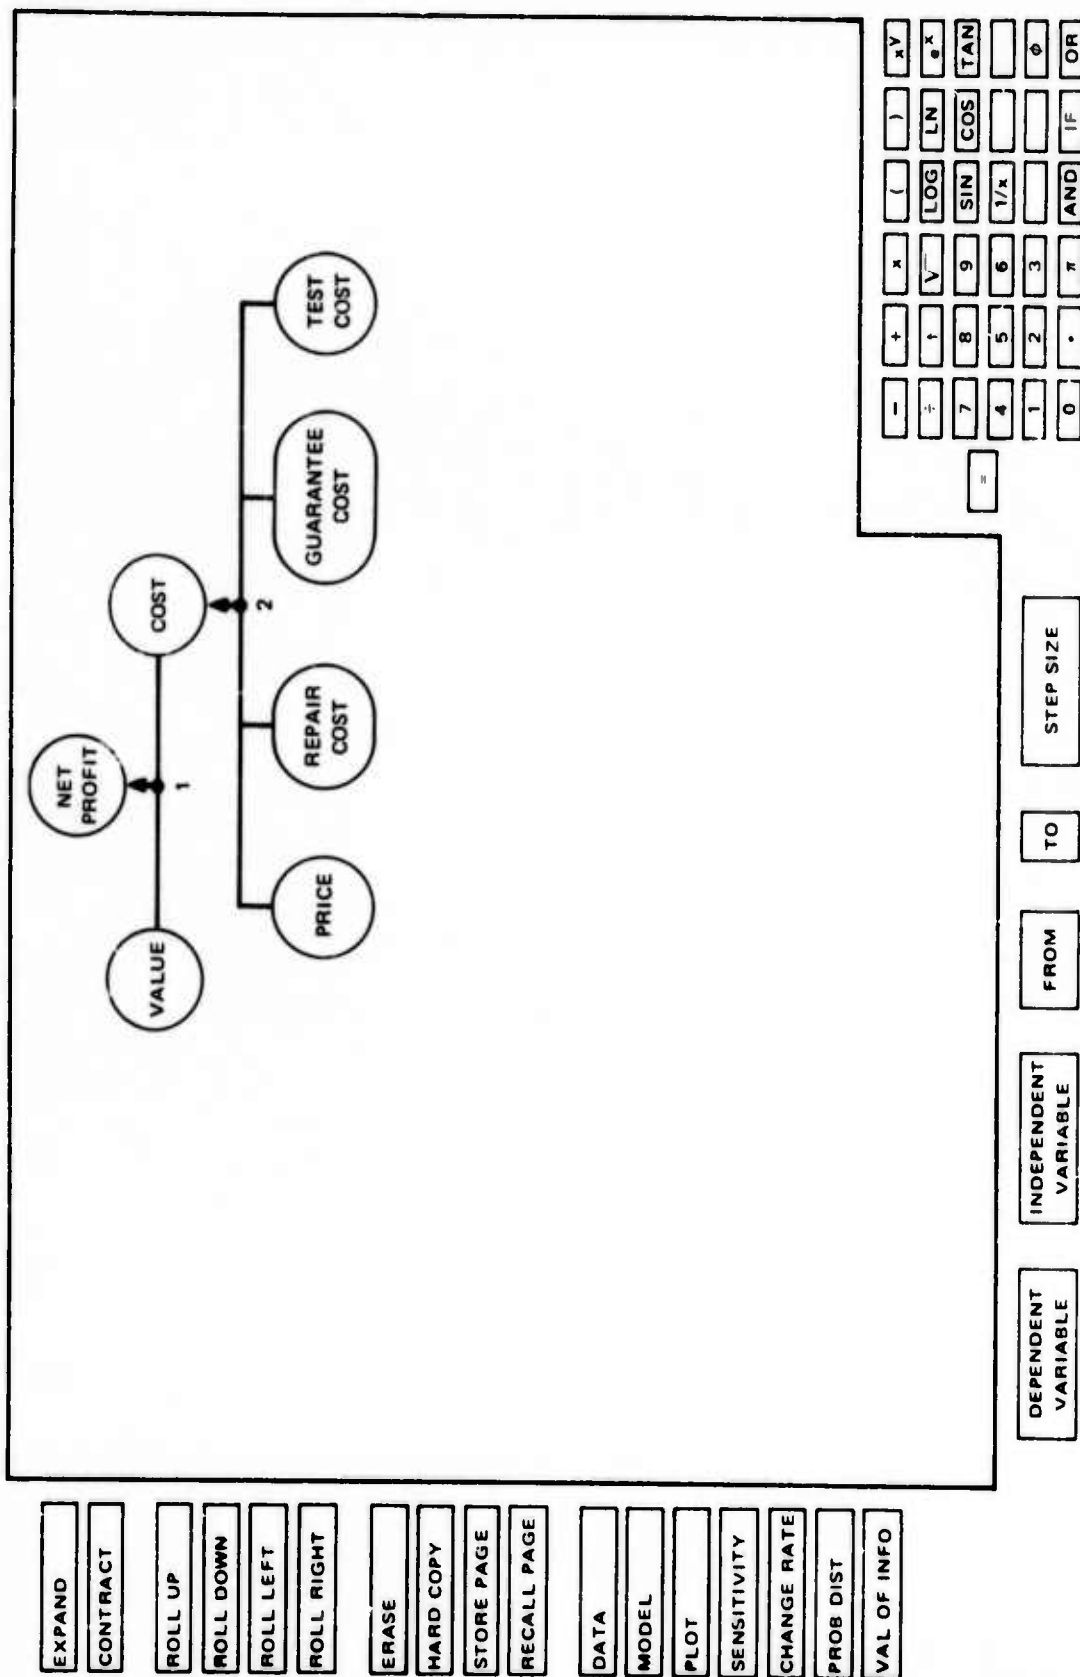
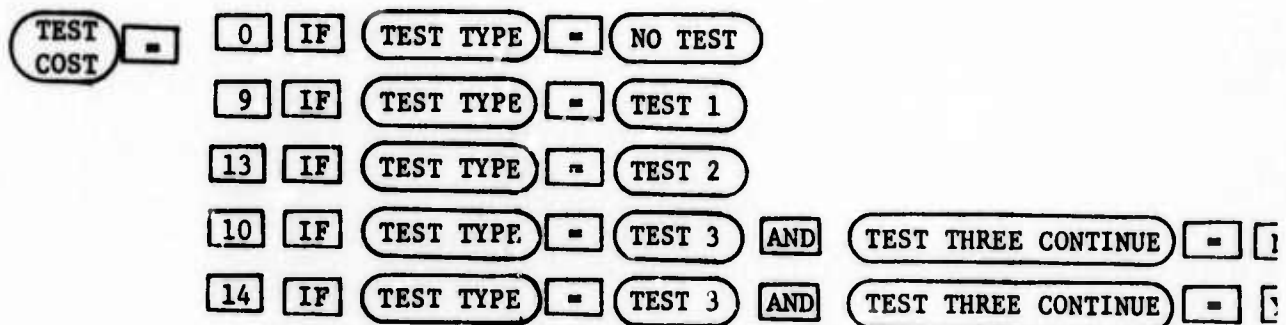


FIGURE 3.15 STEP TWO FOR MODELING THE USED CAR BUYER PROBLEM

Redirecting his attention to the branching structure of Figure 3.15, Joe touches the light pen to TEST COST. The familiar question, this time asking him what he would need to know to define "test cost", appears on the screen. "This is a little more complicated," thinks Joe as he pulls the notes he made of a conversation he had with his mechanic. The dealer placed a time limit of one hour on the car's absence from the lot. In this short time allotment Joe's mechanic said he could at most check only two of the car's ten major systems. In particular he said that he would

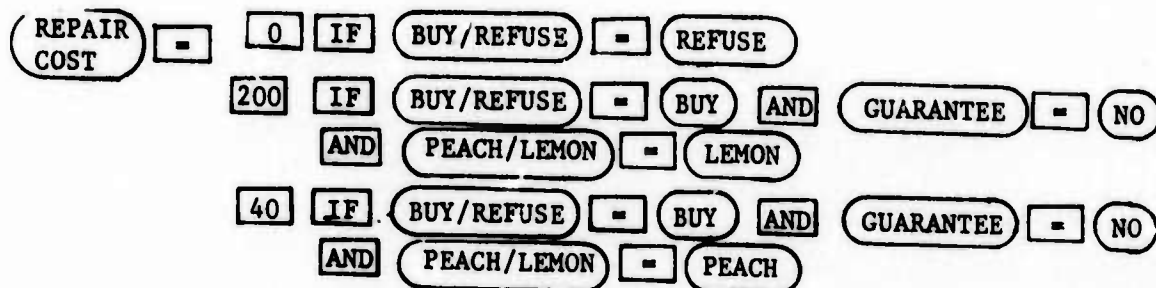
1. test the steering system alone, at a cost of \$9;
2. test the fuel and electrical systems, for a total cost of \$13;
3. test the transmission, at a cost of \$10, report the outcome of the test to Joe, and then proceed to check the differential, at an additional cost of \$4, if he is requested to do so.

"Well, the test cost will depend upon which of the test options I choose," thinks Joe as he types in TEST TYPE. Since "test type" is a decision variable, he immediately adds, DECISION. "And also," he continues, "if I choose the sequential test option, it will depend upon whether I instruct my mechanic to continue and test the differential as well." Joe types TEST THREE CONTINUE. Again, since this is a decision variable, Joe types DECISION. Next, Joe uses the logical "IF" and "AND" buttons on the calculator to write



Joe has defined operator number 3, and the video screen now appears as in Figure 3.16.

Next Joe chooses to define "repair cost." According to the mechanic, it will cost about \$40 to repair a single serious defect in one of the car's major systems, but if 6 defects were to be repaired, the price for all 6 would be only \$200. The anti-lemon guarantee, Joe recalls, cost \$60 and will pay for 50% of the repair costs if the car is a peach, 100% of the costs if it turns out to be a lemon. Joe points the light pen to **REPAIR COST** and then types in PEACH/LEMON, GUARANTEE, and BUY/REFUSE, indicating that repair costs depend upon whether or not the car is a peach or a lemon, whether he purchases the guarantee, and, of course, whether or not he buys the car. Since the entities "buy/refuse" and "guarantee" are decision variables, after each of these he types DECISION. With the calculator for defining the operator, he writes



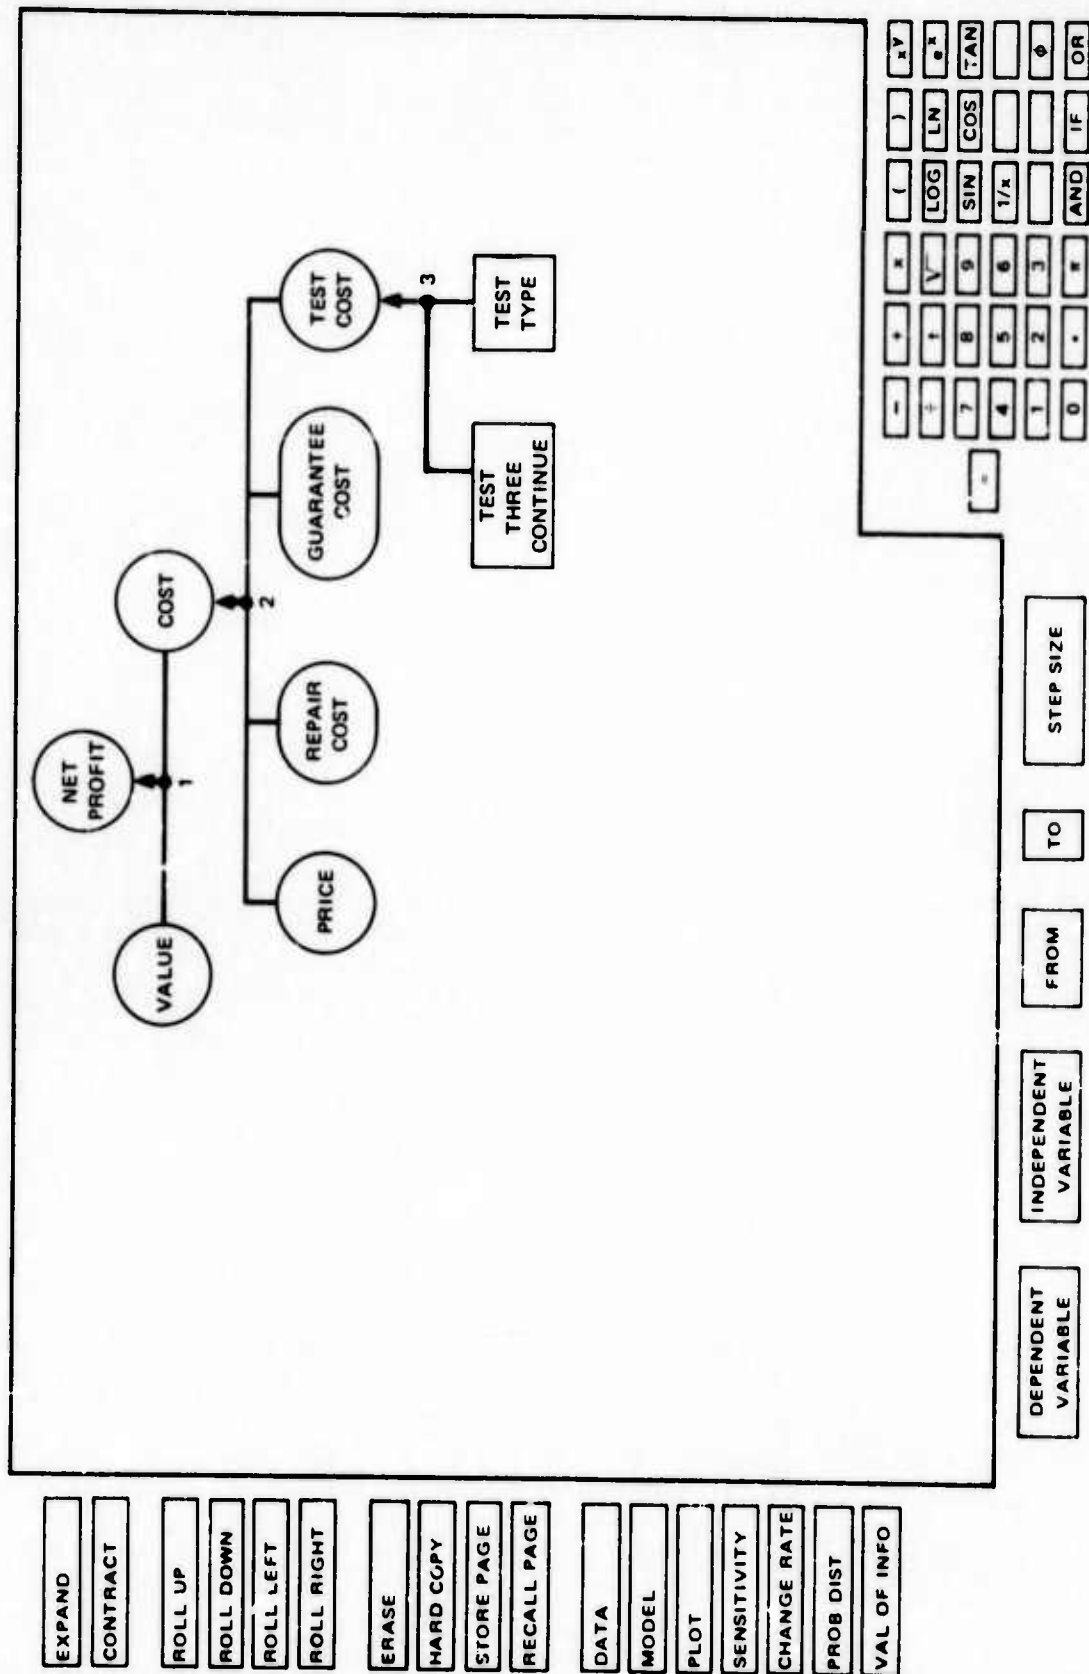
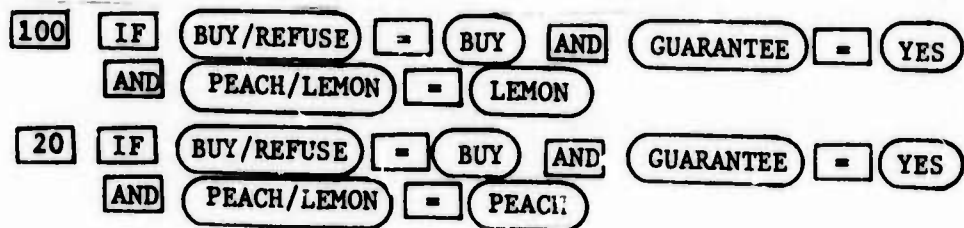
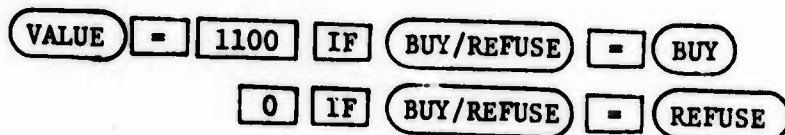


FIGURE 3.16 STEP THREE FOR MODELING THE USED CAR BUYER PROBLEM

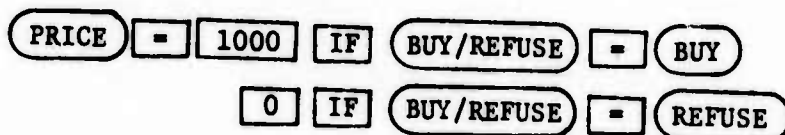


The branching structure now appears as in Figure 3.17.

Joe can see at this point that all of the important system entities appear on his video screen. All that remains now is for him to specify the remaining connections. "Well," he thinks, "'value' depends upon whether or not I buy the car." Joe takes the light pen, pushes **VALUE** and then **BUY/REFUSE**. Operator dot number five appears on the screen and arrows appear connecting the "buy/refuse" block, the new operator dot, and the "value" block. Joe estimates the value of the car to be \$1100, so he specifies operator 5 with



"The asking price for the car is "1,000," Joe continues, "so it will cost me \$1,000 if I decide to buy the car." He touches the light pen first to **PRICE** and then to **BUY/REFUSE**. Operator 6 appears and Joe types



Finally, Joe specifies the dependence of "guarantee cost" on the "guarantee" decision by touching the light pen to the two blocks

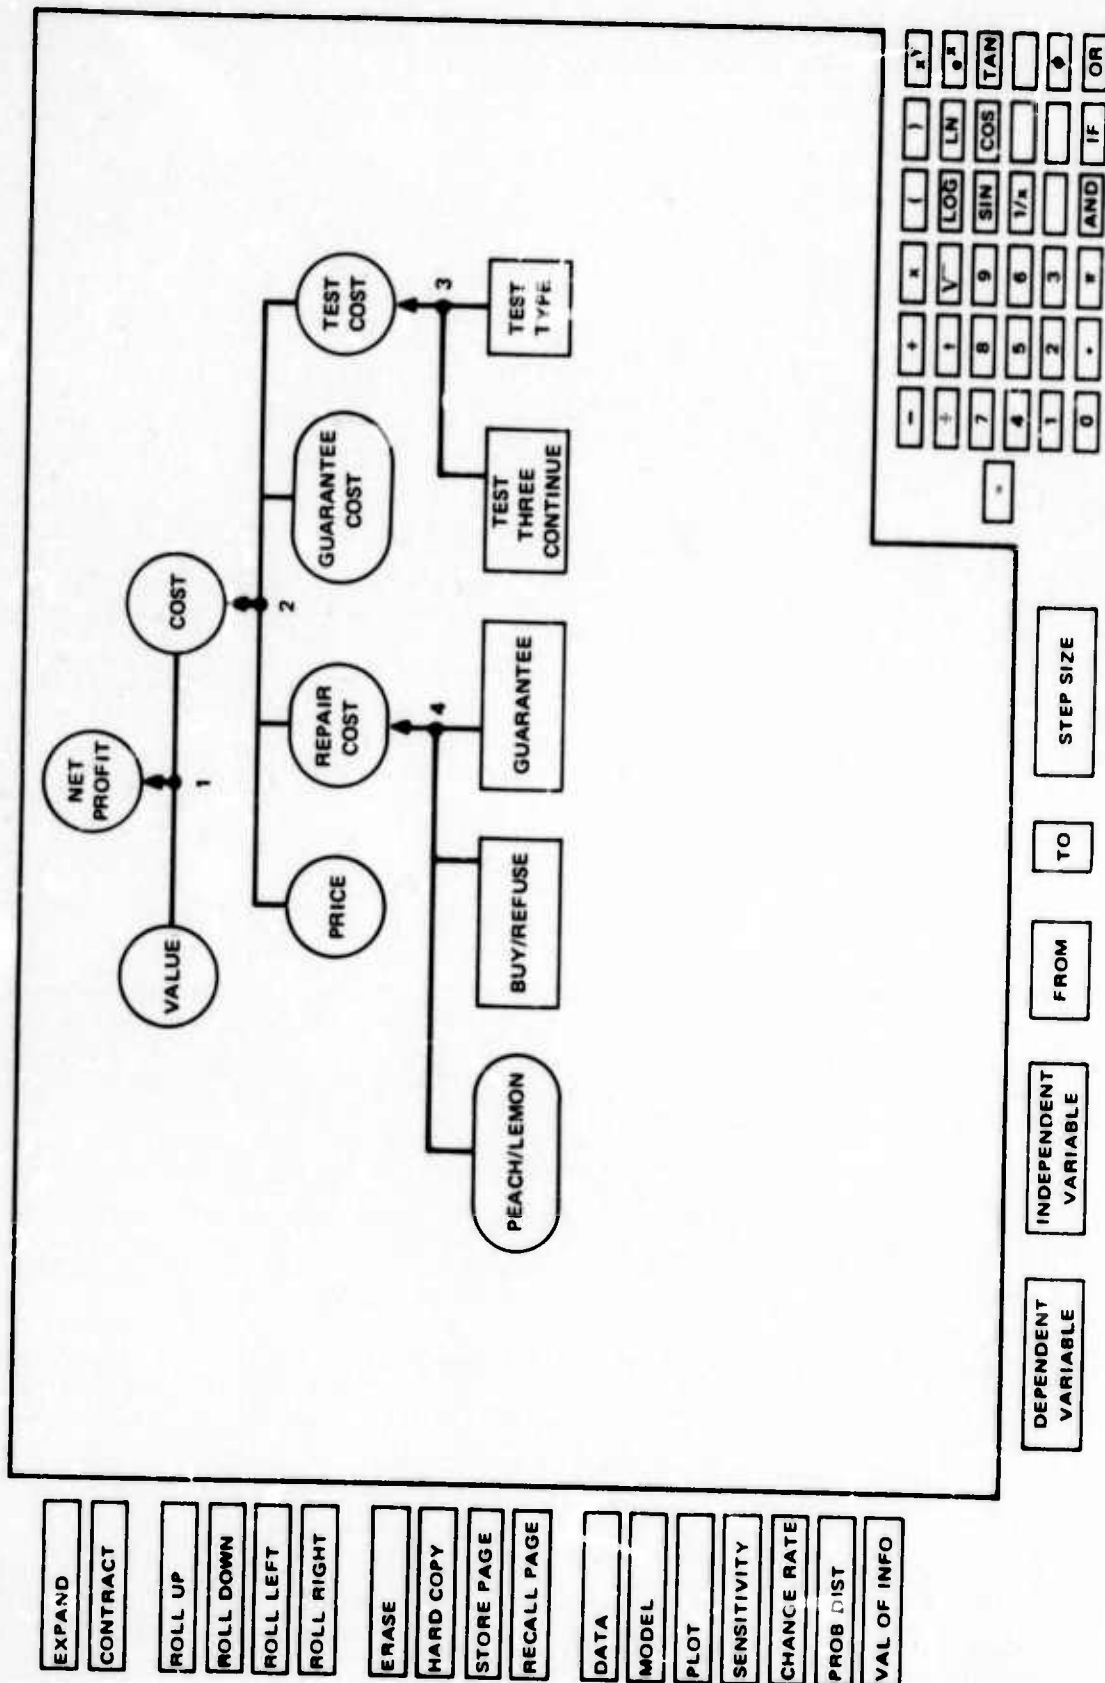
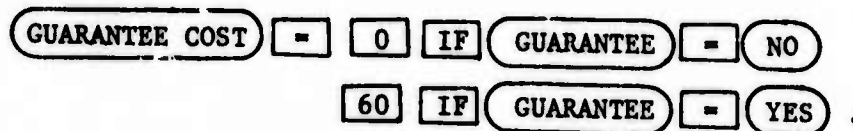


FIGURE 3.17 STEP FOUR FOR MODELING THE USED CAR BUYER PROBLEM

and then defining

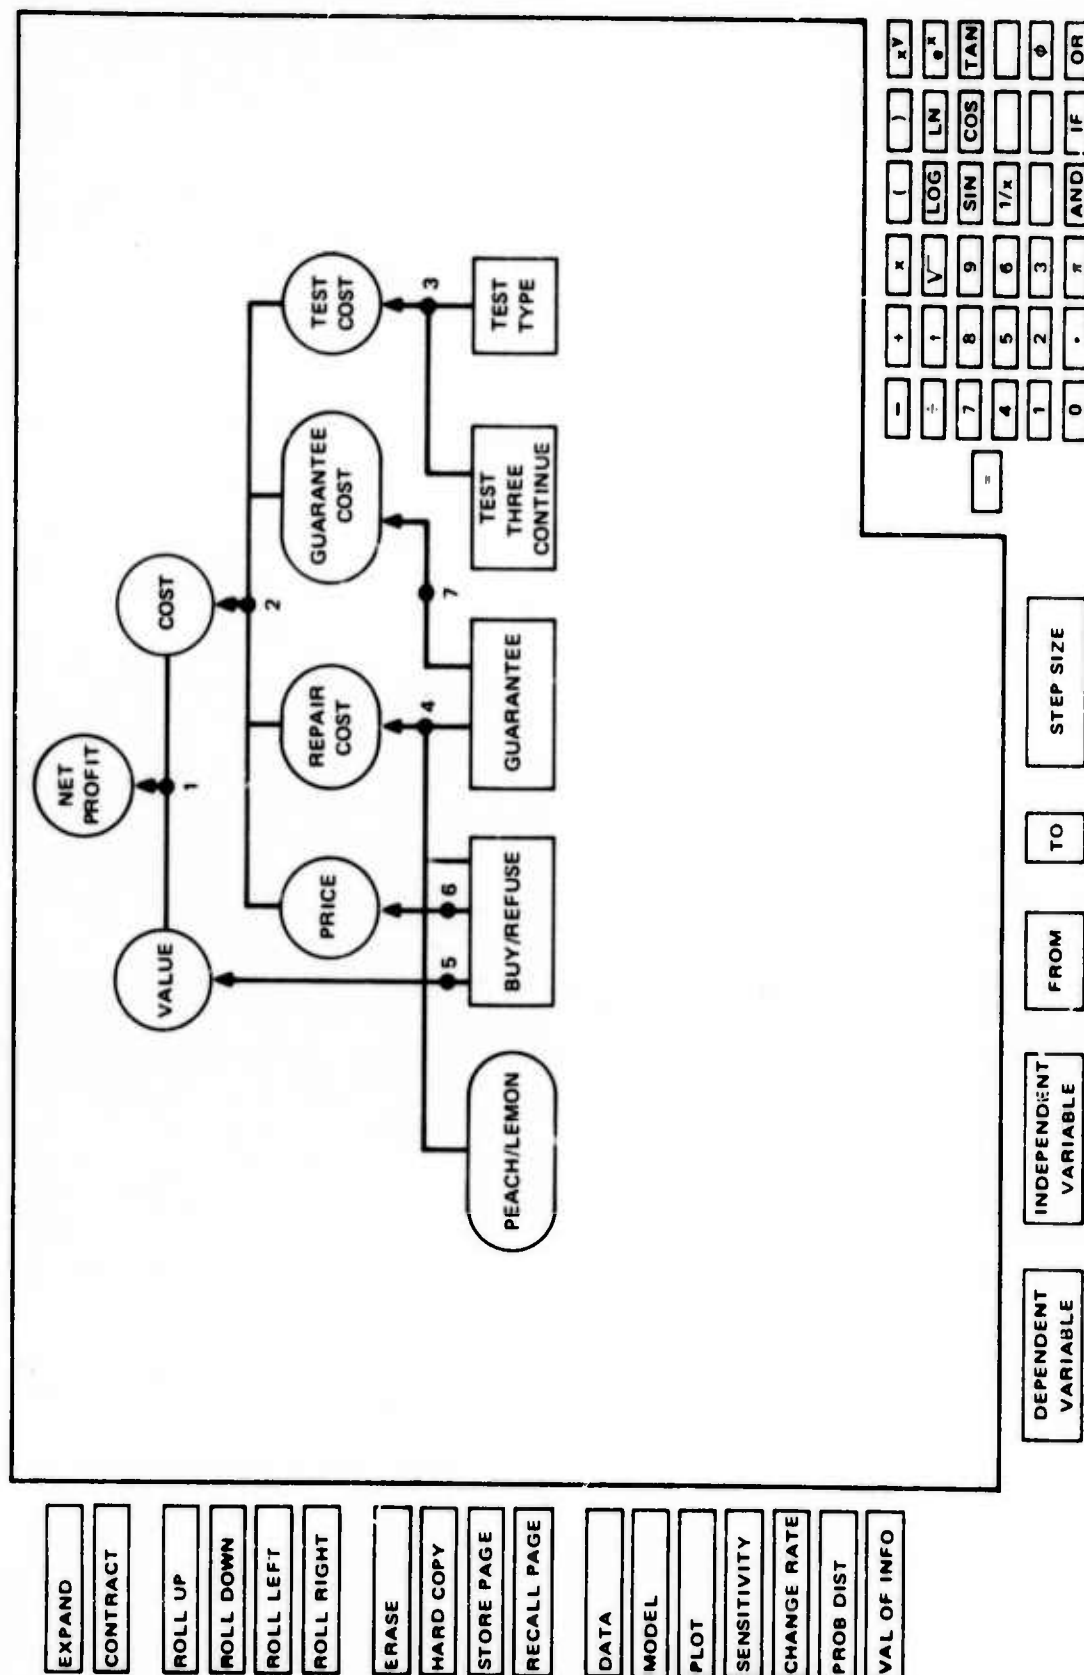


The branching structure appears as shown in Figure 3.18.

3.3.1 Deterministic Analysis

Joe looks closely at the branching structure that now shows on the video screen (Figure 3.18). "That seems like all the connections," he thinks, "but I'd better go over them to make sure." Gradually moving his gaze from the top to the bottom of the figure he observes, "Net profit will equal the value of the decision strategy minus the costs. Value depends upon whether or not I buy the car. Cost equals the sum of price, repair cost, guarantee cost, and test cost. The price I pay depends on whether or not I buy the car. Repair cost depends on whether the car is a peach or a lemon, whether or not I buy the car, and whether or not I guarantee the car. Guarantee cost depends on whether or not I buy the guarantee, and test cost depends on the testing strategy I choose. Yes, it seems to make sense," he thinks. "If I specify those five bottom-most blocks, that should specify net profit."

However, Joe does not feel entirely convinced of his last statement and decides to specify a decision strategy and an outcome for the "peach/lemon" state variable just to see if his model yields a reasonable value for net profit. Joe communicates his intent to CADMUS by touching the control button marked "data". The branching structure of Figure 3.18 vanishes from the video



screen and in its place appears the calculator of Figure 3.19. For each elemental variable (a block in Figure 3.18 having no entering arrows) a corresponding button exists on the face of the calculator. The buttons are categorized according to whether they represent state variables or decision variables. To the right of each state variable button, there is a set of buttons representing the possible outcomes that have been so far defined for that variable. Similarly, to the right of each decision variable button there is a set of buttons representing possible alternatives for that decision variable. Joe touches the light pen to the buttons as he would on a conventional calculator so as to write the following equations:

PEACH/LEMON	=	LEMON
BUY/REFUSE	=	BUY
GUARANTEE	=	NO
TEST TYPE	=	NO TEST
TEST THREE CONTINUE	=	NO

He then touches **CALCULATE** . The screen goes blank and an instant later shows the message

NET PROFIT = -100.

Joe mentally verifies this result. "If I accept the deal," he reasons, "I will pay \$1,000 for a car I believe to be worth \$1,100. However, if the car turns out to be a lemon, I'll have to pay \$200 in repair bills. Therefore, I would end up losing \$100 dollars in

EXPAND	DECISION VARIABLES	ALTERNATIVES	STATE VARIABLES	OUTCOMES
CONTRACT	BUY/REFUSE	BUY	PEACH/LEMON	PEACH
ROLL UP		REFUSE		LEMON
ROLL DOWN				
ROLL LEFT	GUARANTEE	YES		
ROLL RIGHT		NO		
ERASE				
HARD COPY				
STORE PAGE		NO TEST		
RECALL PAGE	TEST TYPE	TEST 1		
		TEST 2		
DATA		TEST 3		
MODEL				
PLOT		YES		
SENSITIVITY	TEST THREE CONTINUE	NO		
CHANGE RATE				
PROB DIST				
VAL OF INFO				

DEPENDENT VARIABLE	INDEPENDENT VARIABLE	FROM	TO	STEP SIZE

CALCULATE	=	-	+	*	/	()	x ^y
		7	8	9	√	LOG	LN	e ^x
		4	5	6	1/x	SIN	COS	TAN
		1	2	3				φ
		0	.	π	AND	IF		OR

FIGURE 3.19 CONTROL PANEL FOR DETERMINISTIC ANALYSIS

the deal." Feeling encouraged, Joe decides to check his model with a few more calculations. Typing ANALYSIS, the calculator of Figure 3.19 again appears. Joe touches the buttons so as to write

= ,
 =
 =
 =
 =

Joe pushes and the result appears on the screen.

NET PROFIT = 7, 27

Joe is somewhat surprised to see that his model indicates that he will make more money, \$27, if the car turns out to be a lemon, than the \$7 he will make if it turns out to be a peach. He verifies the result as follows: "If I buy the car, I make \$100 since the price is that much below what I believe to be the value of that make, model, and year. Testing decision number two will cost me \$13 and the guarantee costs \$60. If the car is a peach, the guarantee will pay for 50% of the \$40 repair costs. That means my net profit would be $\$100 - \$13 - \$60 - \$20 = \$7$. If the car is a lemon, the guarantee will pay for all repair costs so the net profit will be $\$100 - \$13 - \$60 = \27 . That checks," thinks Joe.

3.3.2 Influence Specification

The branching structure of Figure 3.18 lacks an important specification of a decision problem, probabilistic dependencies

representing the effects of information. Such dependencies may be of crucial importance for decision making and cannot be ignored. In the Used Car Buyer Problem, for example, the outcome of the test that might be performed by the mechanic is not included as an entity in the model of Figure 3.18 because that outcome does not have any direct effect on net profit. On the other hand, the outcome of such a test will influence Joe's estimation of the likelihood that the car is a peach. Since Joe's purchasing decision will be strongly dependent on his perception of the odds that the car is a peach, the outcome of the test is important information.

Specifying Information Flow With CADMUS

One method for describing informational flow with CADMUS is to incorporate an influence diagram (see Section 2.4.2) into the deterministic branching structure representing the problem's value model. Let us suppose that our used car buyer, Joe, chooses to use this technique. Joe informs CADMUS by typing INFLUENCE DIAGRAM on the terminal. CADMUS answers:

IF YOU WISH YOUR DECISION STRATEGY TO BE CONDITIONED UPON INFORMATION OR STATE VARIABLE OUTCOMES WHICH MIGHT BECOME KNOWN DURING THE DECISION PROCESS, LIST ALL SUCH INFORMATION OR OUTCOME VARIABLES WHICH HAVE NOT YET BEEN INCLUDED IN THE BRANCHING STRUCTURE.

The information generated in the course of Joe's decision process will be the outcomes of any tests that he asks his mechanic to perform. Therefore, Joe types

TEST OUTCOME

However, if Joe chooses "test 3", the sequential testing strategy, Joe has the option to ask that a second system in the car be tested after he learns the result of the first test. Hence, the outcome of the decision to continue testing is a possible piece of information upon which Joe will wish to base his purchase and guarantee decisions. Joe types

CONTINUE TEST OUTCOME

and signals CADMUS that this is the last information variable to be added to the model. The branching structure appears on the screen together with entity blocks representing the added information variables (Figure 3.20).

Joe must now interconnect the added information variables and the elemental state and decision variables in the form of an influence diagram. In many problems this connection is facilitated by thinking first of the order in which the decisions must be made and the state variable outcomes become known. Joe must first decide on the "test type." Then the "test outcome" becomes known, and Joe may have the option to ask that another test be performed ("test three continue"). If he has chosen this option the "continue test outcome" is revealed and Joe must decide to buy or refuse the car deal ("buy/refuse"), and, if he buys the car, whether or not to have it guaranteed ("guarantee"). Only after he makes these decisions does he finally learn for sure whether the car is a peach or a lemon ("peach/lemon").

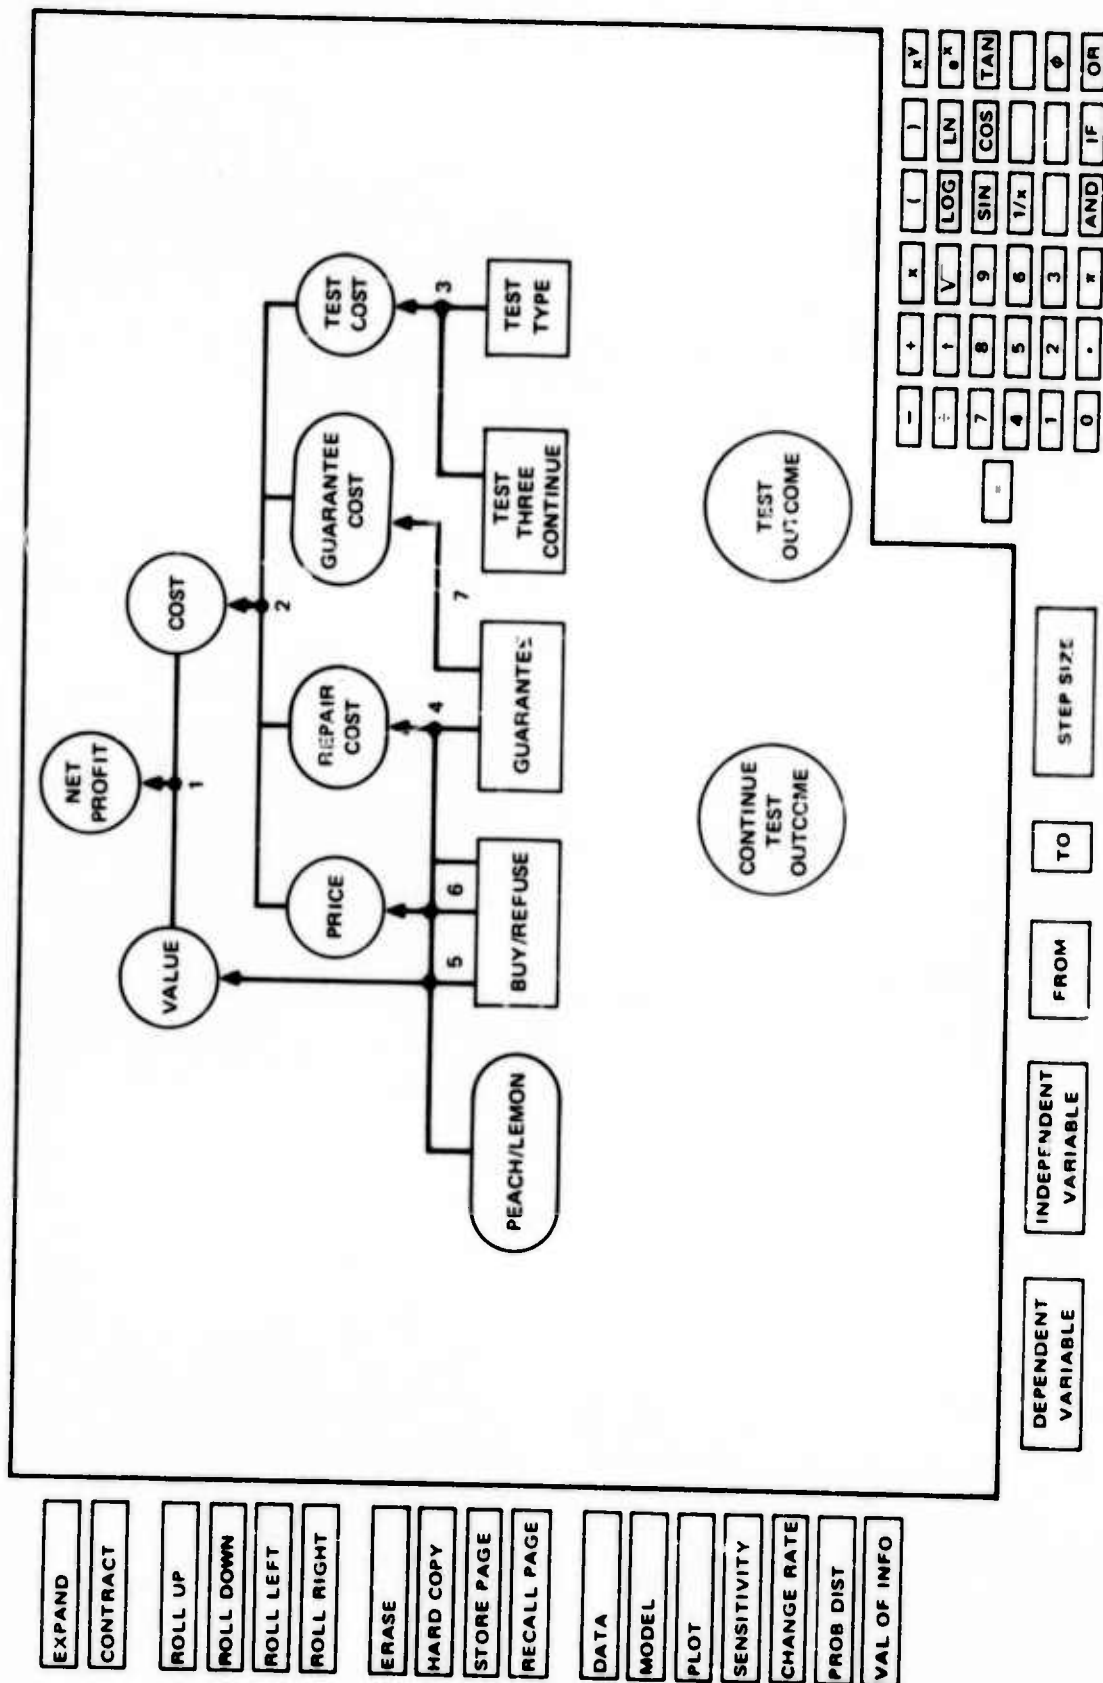


FIGURE 3.20 USED CAR BUYER VALUE STRUCTURE PLUS INFORMATION OUTCOMES

The probability of the car being a peach or a lemon may depend on the test outcomes. Joe must represent this probabilistic dependence by drawing influence arrows from the "continue test outcome" block and from the "test outcome" block to the "peach/lemon" block. Influence arrows may be drawn using CADMUS by touching the light pen first to the entity block from which the arrow is to originate and then to the block at which it is to terminate. An arrow connecting the two blocks will immediately materialize on the video screen.

The "buy/refuse" decision as well as the "guarantee" decision ought to depend upon the testing outcomes so influence arrows are drawn from each of the test outcome blocks to the "buy/refuse" and "guarantee" decision blocks. A "continue test outcome" can only exist if the decision is made to continue testing, so an influence arrow must be drawn from the "test three continue" block to the "continue test outcome" block. An arrow is drawn also from the "test outcome" block to the "continue test outcome" block because the assessment of the probabilities of the various outcomes for the second test will most likely be influenced by the outcome of the first test. Since the "test three continue" decision may only be made if the third test type is chosen, and since this decision may be made with knowledge of the first "test outcome", arrows are drawn connecting the test type and test outcome blocks to the test three continue block. The complete branching structure for the Used Car Buyer Problem, including the appended influence diagram, is now on the video screen. This is

illustrated in Figure 3.21.

3.3.3 Constraint Statements

One thing, however, bothers Joe. He has not made it clear yet that some state variable outcomes and some decision alternatives in his problem exist only under certain conditions. For example, the "test three continue" decision cannot be made unless the testing strategy is chosen to be "test 3". To make such problem constraints specific, Joe types "constraints" and the calculator of Figure 3.22 appears on the screen. Joe pushes the calculator buttons so as to write

NO TEST THREE CONTINUE IF TEST TYPE NOT = TEST 3
NO CONTINUE TEST OUTCOME IF TEST THREE CONTINUE = NO
NO TEST OUTCOME IF TEST TYPE = NO TEST
NO GUARANTEE IF BUY/REFUSE = REFUSE .

3.3.4 Probability Specification

To calculate the optimal decision strategy under uncertainty, probability assignments must be expressed in a form which is appropriate to the given problem structure. Frequently, this will not be the form that the decision maker finds most convenient for assessing probabilities. Consequently, calculations are usually required to convert from one form to another. This conversion is accomplished using Bayes' rule.

CADMUS will perform Bayes' rule automatically so the decision maker may specify probabilities under any conditioning assumptions he desires. A number of methods to aid the decision maker in

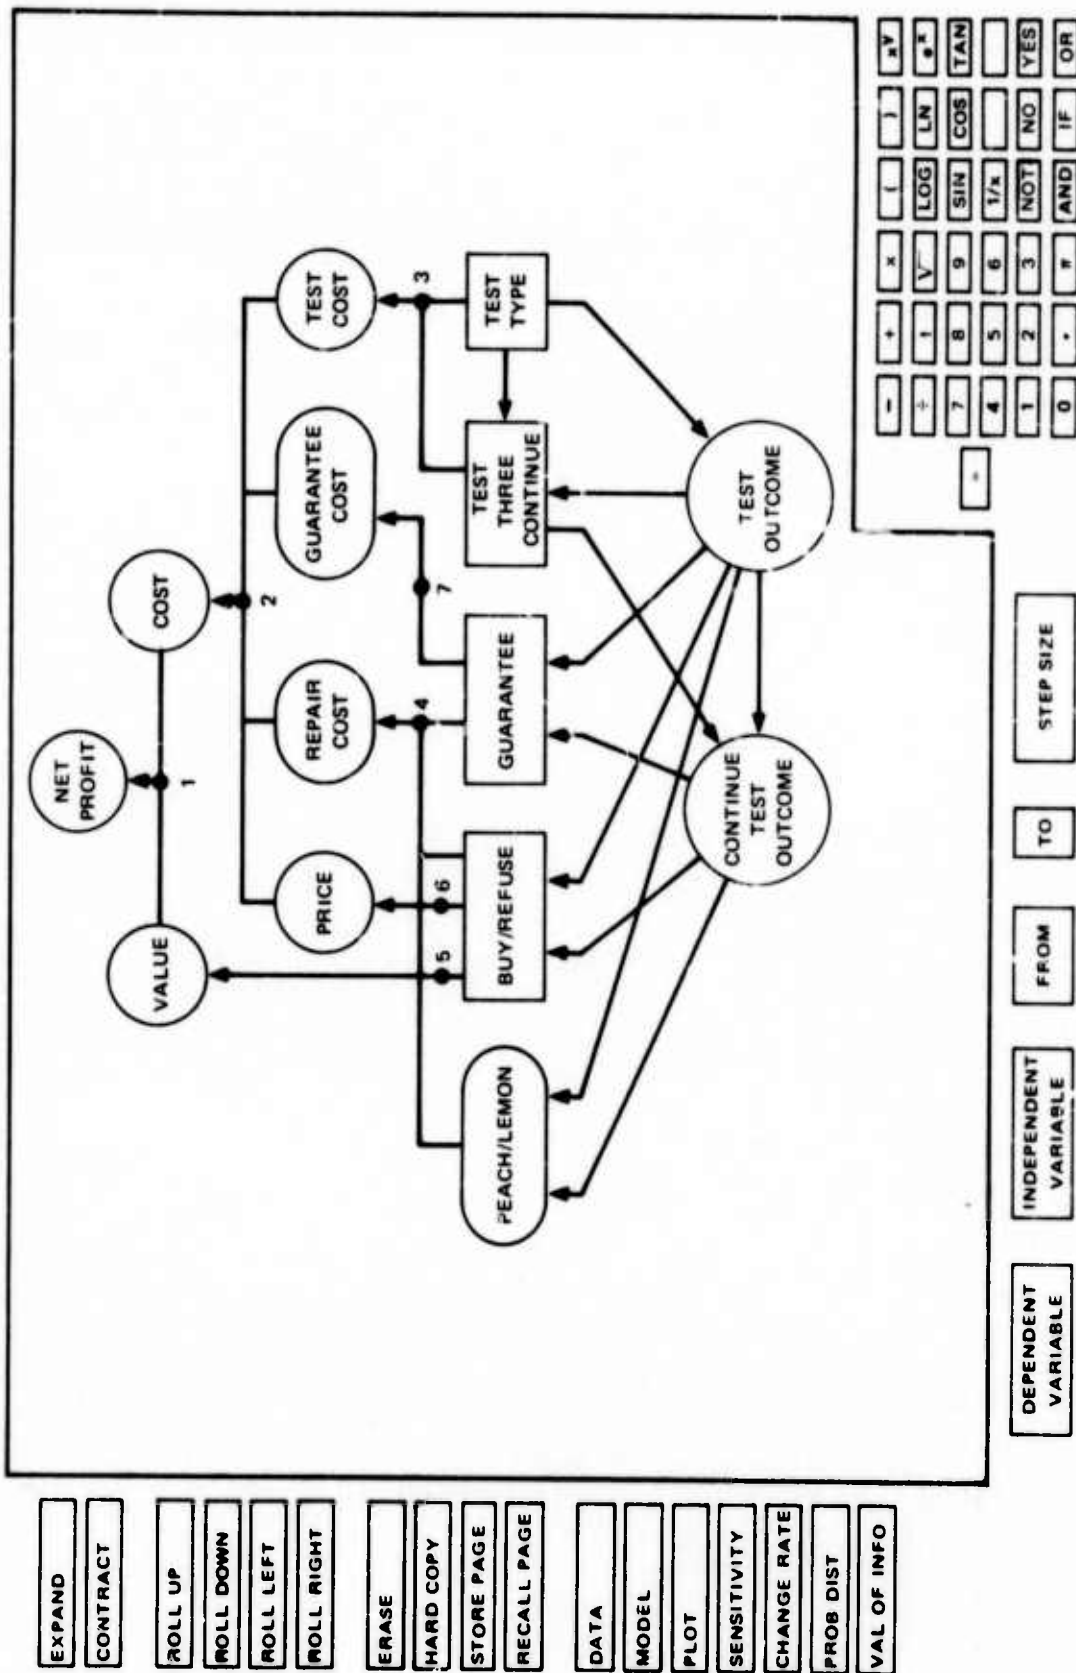


FIGURE 3.21 USED CAR BUYER VALUE STRUCTURE WITH APPENDED INFLUENCE DIAGRAM

EXPAND

CONTRACT

ROLL UP

ROLL DOWN

ROLL LEFT

ROLL RIGHT

ERASE

HARD COPY

STORE PAGE

RECALL PAGE

DATA

MODEL

PLOT

SENSITIVITY

CHANGE RATE

PROB DIST

VAL OF INFO

DECISION VARIABLES

BUY/REFUSE

GUARANTEE

TEST TYPE

TEST THREE CONTINUE

ALTERNATIVES

BUY

REFUSE

YES

NO

NO TEST

TEST 1

TEST 2

TEST 3

YES

NO

STATE VARIABLES

PEACH/LEMON

TEST OUTCOME

CONTINUE TEST OUTCOME

OUTCOMES

PEACH

LEMON

DEPENDENT VARIABLE

INDEPENDENT VARIABLE

FROM

TO

STEP SIZE

-

+

7

4

1

0

1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

44

45

46

47

48

49

50

51

52

53

54

55

56

57

58

59

60

61

62

63

64

65

66

67

68

69

70

71

72

73

74

75

76

77

78

79

80

81

82

83

84

85

86

87

88

89

90

91

92

93

94

95

96

97

98

99

100

FIGURE 3.22 CALCULATOR FOR CONSTRAINT SPECIFICATION

assessing probabilities for the uncertain state variables are provided by CADMUS.

For the Used Car Buyer Problem, probability assessment is most natural in the order represented in Figure 3.23. Joe believes that there is a one in five chance that the used car he is interested in is a lemon. So, the a priori probability of the car being a lemon is 0.2, and the a priori probability of it being a peach is 0.8. Joe must now determine the probabilities of the various possible test outcomes given that the car is a lemon or given that it is a peach. This is a simple problem in combinatorial mathematics.

In a lemon exactly six of the car's ten systems will be defective. Suppose that one major system of the car is tested. Since the car is a lemon, there is probability 6 in 10, or 0.6, that the system being tested happens to be one of the defective systems. Therefore, if the car is a lemon and the "test type" is "test 1", then there is probability 0.6 that the test outcome will be "defect" and a probability 0.4 that this outcome will be "no defect".

Suppose that a defective system is found on the first test. There are nine remaining untested systems in the car, and, out of those nine, five are defective. Therefore, if a second system is tested there is probability 5 in 9, or $5/9$, that this system will likewise be defective. Hence, the probability of finding two defective systems in two tests is $0.6 \times 5/9 = 1/3$. Stated differently, if the car is a lemon, and if the "test type" is "test 2", then there is probability $1/3$ that the test outcome is "two defects".

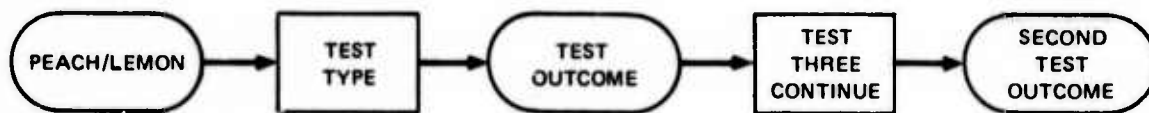


FIGURE 3.23 NATURAL CONDITIONING ORDER FOR PROBABILITY SPECIFICATIONS FOR THE USED CAR BUYER PROBLEM

By continuing with such reasoning, Joe is able to determine all the probabilities necessary for the specification of the informational model illustrated in Figure 3.23.

Joe must now input these probabilities to CADMUS. By typing PROBABILITY SPECIFICATION on the terminal a probability specification "calculator" is made to appear on the screen (Figure 3.24). This calculator will minimize the work needed to input state variable probabilities. Additional "buttons" may be defined as needed on the calculator to correspond to state variable outcomes which have not yet been defined. The decision maker merely touches the light pen to the screen and then types in the outcome name on the terminal keyboard. In this way Joe obtains buttons on his calculator for the various possible testing outcomes.

The logical and mathematical expressions that Joe enters with the calculator are shown in Figure 3.25. Joe may ask CADMUS to illustrate his probability specifications in a probability tree. If he does so the video screen would appear as shown in Figure 3.26.

3.3.5 Representation of the Decision in a Decision Tree

To check his model Joe may ask to see it illustrated in decision tree form. If at this point Joe were to type DECISION TREE, the tree structure of Figure 3.27 would appear on the video screen. (Ignore for the moment the arrows shown on some decision branches and the figures within the circles.) The reader may verify that Joe has completely specified his model. Joe's first decision is

EXPAND

CONTRACT

ROLL UP

ROLL DOWN

ROLL LEFT

ROLL RIGHT

ERASE

HARD COPY

STORE PAGE

RECALL PAGE

DATA

MODEL

PLOT

SENSITIVITY

CHANGE RATE

PRDB DIST

VAL OF INFO

DECISION VARIABLES

BUY/REFUSE

GUARANTEE

TEST TYPE

TEST THREE CONTINUE

ALTERNATIVES

BUY

REFUSE

YES

NO

NO TEST

TEST 1

TEST 2

TEST 3

YES

NO

STATE VARIABLES

PEACH/LEMON

TEST OUTCOME

CONTINUE TEST OUTCOME

OUTCOMES

PEACH

LEMON

DEFECT

NO DEFECT

DEFECT

NO DEFECT

IF PEACH/LEMON = PEACH
IF TEST TYPE = TEST 1
TEST OUTCOME = DEFECT WITH PROBABILITY 0.1
NO DEFECT WITH PROBABILITY 0.9

DEPENDENT VARIABLE

INDEPENDENT VARIABLE

FROM

TD

STEP SIZE

-

+

*

/

7

8

9

4

1

0

=

(

)

x

y

LN

LOG

SIN

COS

TAN

1/x

3

π

AND

IF

DR

FIGURE 3.24 CALCULATOR FOR SPECIFYING PROBABILITIES

IF PEACH/LEMON = PEACH

IF TEST TYPE = TEST 1 OR TEST 3

TEST OUTCOME = $\left\{ \begin{array}{ll} \text{DEFECT} & .1 \\ \text{NO DEFECT} & .9 \end{array} \right.$

IF TEST TYPE = TEST 2

TEST OUTCOME = $\left\{ \begin{array}{ll} \text{NO DEFECT} & 4/5 \\ 1 \text{ DEFECT} & 1/5 \\ 2 \text{ DEFECT} & 0 \end{array} \right.$

IF TEST TYPE = TEST 3

IF TEST OUTCOME = DEFECT

SECOND TEST OUTCOME = $\left\{ \begin{array}{ll} \text{DEFECT} & 0 \\ \text{NO DEFECT} & 1 \end{array} \right.$

IF TEST OUTCOME = NO DEFECT

SECOND TEST OUTCOME = $\left\{ \begin{array}{ll} \text{DEFECT} & 1/9 \\ \text{NO DEFECT} & 8/9 \end{array} \right.$

Figure 3.25: PROBABILITY SPECIFICATION FOR THE
USED CAR BUYER PROBLEM

IF PEACH/LEMON = LEMON

IF TEST TYPE = TEST 1 OR TEST 3

TEST OUTCOME =	DEFECT	.6
	NO DEFECT	.4

IF TEST TYPE = TEST 2

TEST OUTCOME =	NO DEFECT	2/15
	1 DEFECT	8/15
	2 DEFECT	1/3

IF TEST TYPE = TEST 3

IF TEST OUTCOME = DEFECT

SECOND TEST OUTCOME =	DEFECT	5/9
	NO DEFECT	4/9

IF TEST OUTCOME = NO DEFECT

SECOND TEST OUTCOME =	DEFECT	2/3
	NO DEFECT	1/3

Figure 3.25: PROBABILITY SPECIFICATION FOR THE
USED CAR BUYER PROBLEM

(Concluded)

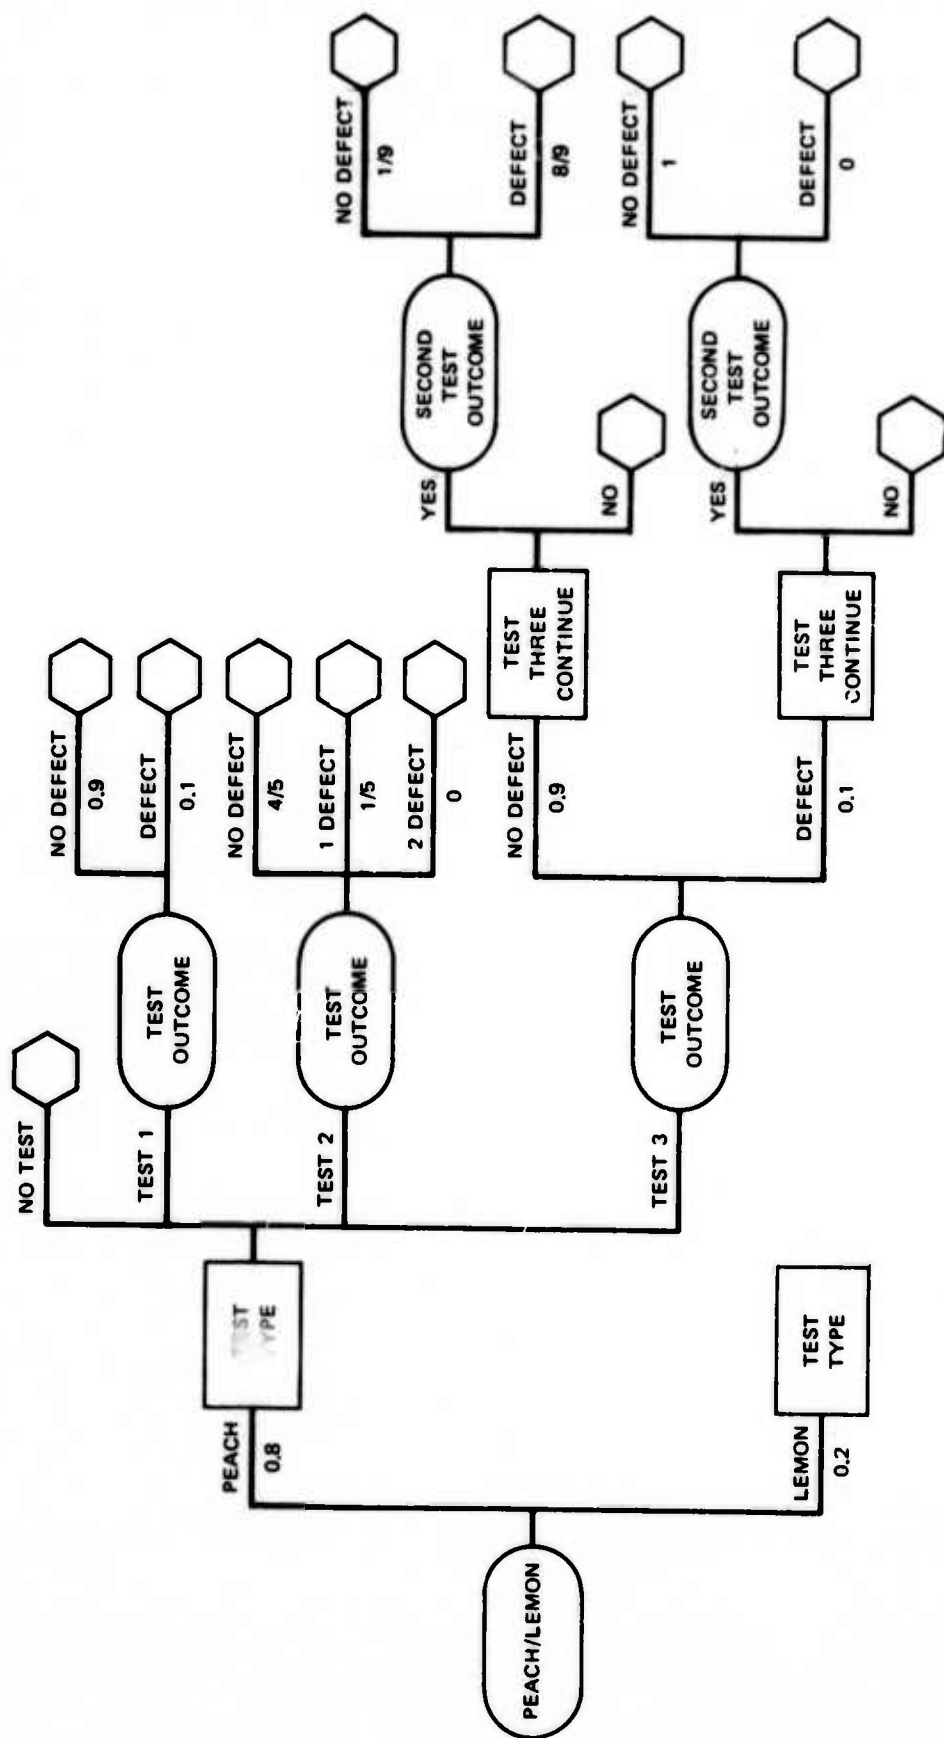


FIGURE 3.26 TREE ILLUSTRATING JOE'S PROBABILITY SPECIFICATIONS

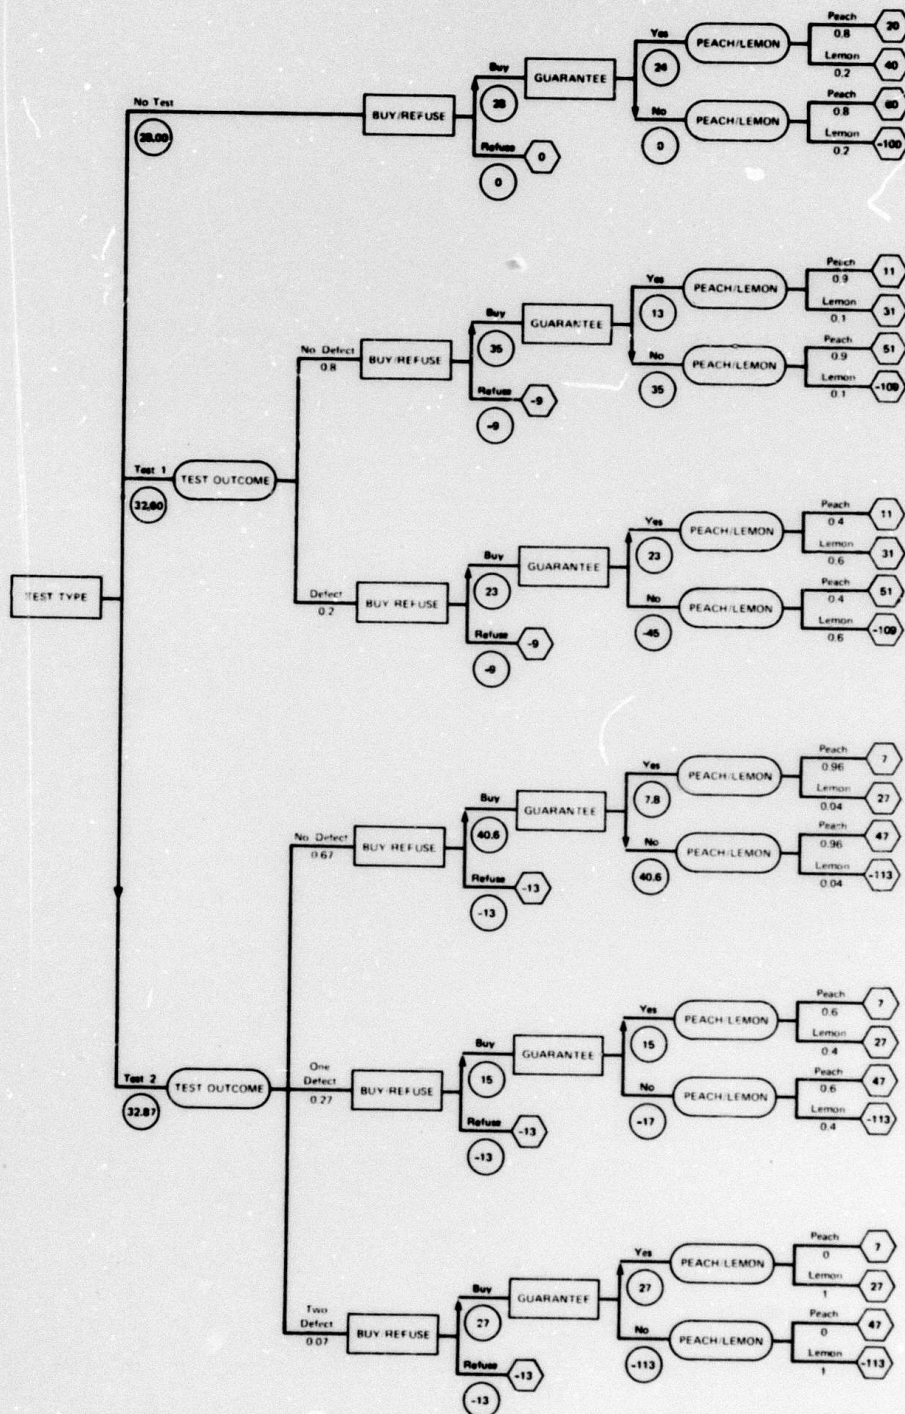


FIGURE 3.27 DECISION TREE FOR THE USED CAR BUYER PROBLEM

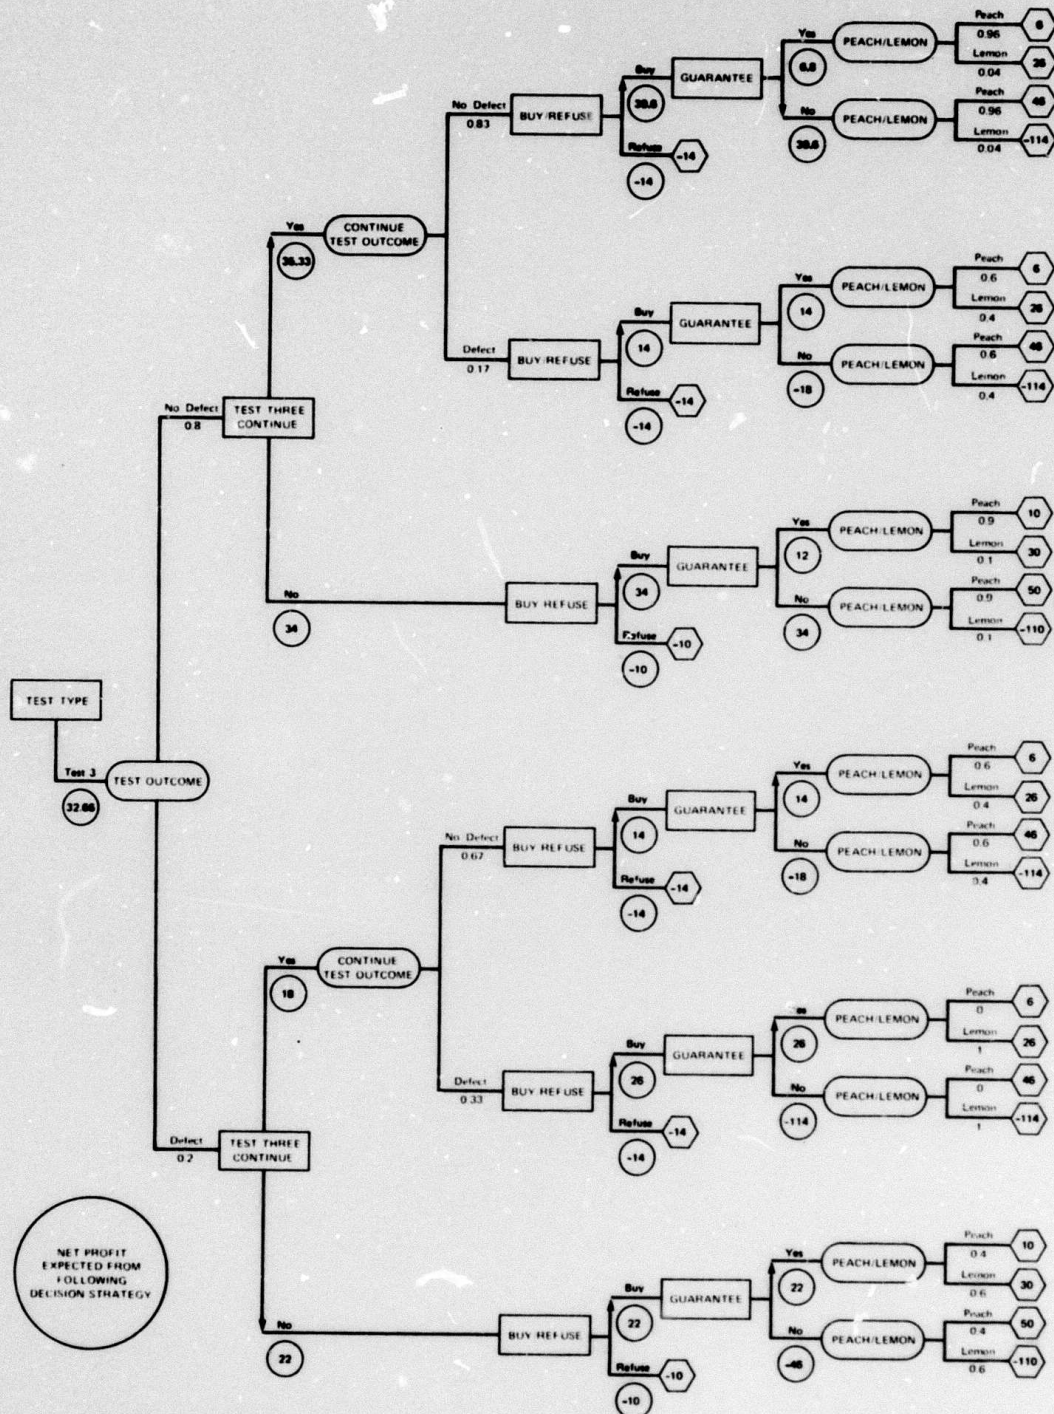


FIGURE 3.27 DECISION TREE FOR THE USED CAR BUYER PROBLEM (Concluded)

to choose whether or not to have the car checked and if so, which of the three testing strategies to choose. If, for example, he chooses "test type" "test 2", two of the car's 10 systems will be checked and Joe will learn that either no defects, one defect, or two defects were found. The probabilities of these various outcomes have been calculated by CADMUS using Bayes' rule to be 0.67, 0.27, and 0.07. After learning which outcome actually occurs, Joe must decide to buy or refuse the car and if he buys it, whether or not to purchase the guarantee. If he buys the car, he will finally learn whether the car is a peach or a lemon. Joe has now completely specified his model.

The complexity of the decision problem is made apparent in the decision tree structure of Figure 3.27. Decision makers more experienced with decision trees would have the option of inputting their model structure directly in the form of a decision tree. This could be accomplished through CADMUS by drawing the tree structure with the light pen and then typing the parameters associated with the tree's nodes and branches through the terminal.

If Joe were to push the "optimize" control button, CADMUS would calculate the expected value associated with each alternative for each decision in the problem. The results of such calculations are shown enclosed within circles under each decision branch of the tree in Figure 3.27. At this point CADMUS is able to identify the decision strategy that maximizes expected net profit. CADMUS may be requested to read out this strategy as a set of

logical "if-then" type statements or it may graphically represent the optimal expected value strategy by placing arrows on the preferred decision alternatives on the decision tree as illustrated in Figure 3.27. We see that Joe will maximize his expected net profit if he instructs his mechanic to perform "test 2", that is test the fuel and electrical systems, and then buys the anti-lemon guarantee only if a defect is found. This strategy will produce an expected net profit of \$32.87.

3.3.6 Utility Specification

Tests are supplied by CADMUS to enable the decision maker to determine whether it is necessary to express his risk attitude as a utility function. If significant risk aversion (or risk preference) is indicated, CADMUS will supply procedures for eliciting an appropriate utility function.

3.4 CONCLUSIONS AND SUMMARY

A branching structure has been described which allows a decision maker to express compactly the structure of a decision problem. Although the technique was originally developed to describe single stage continuous deterministic models, we have seen that the method may be expanded so as to be used in sequential problems with discrete events and significant uncertainties. The system described is an illustration of how a computerized graphical structure may help a decision maker to formulate a decision problem. Current research has demonstrated the feasibility of such a system but additional research will be required to determine exact software specifications.

REFERENCES

- [1] Baumol, W. J. Economic Dynamics, 2nd edn., The Macmillan Company, New York, 1959, pp. 92-93.
- [2] Dahl, O. J., and Nygaard, K., SIMULA - A Language for Programming and Description of Discrete Event System Introduction and User's Manual. Norwegian Computer Centers, 1967.
- [3] Groner, G.F., et al., "BIOMOD: An Interactive Computer Graphics System for Modeling", The Rand Corporation, R-617-NIH, July 1971.
- [4a] Howard, R.A., "The Used Car Buyer", Reprint, Stanford Research Institute, Menlo Park, California, 1974.
- [4b] Howard, R. A., EES 221 Course Notes, Department of Engineering-Economic Systems, Stanford University, Stanford, California, 1974.
- [4c] Howard, R. A., "Proximal Decision Analysis," Management Science, Vol. 17, No. 9, May 1971.
- [5] International Business Machines Corporation, "IBM General Purpose Simulation Systems User's Manual," Form SH20-0851, 1970.
- [6] International Business Machines Corporation "System/360 Continuous Systems Modeling Program (360A-CX-16X) Application Description," Form H-20-240-2, August 1968.
- [7] Kiviat, P.J., Villaneuva, P., and Markowitz, H.M., The SIMSCRIPT II Programming Language. Englewood Cliffs, N.J.: Prentice-Hall, 1968.
- [8] Marschak, Thomas and Richard Nelson, "Flexibility, Uncertainty, and Economic Theory," *Metroeconomica*, April-Dec., 1962 pp. 42-58.
- [9] Merkhofer, M.W., "Flexibility and Decision Analysis," PhD. Dissertation, Stanford University, to be published, 1975.

REFERENCES (Continued)

- [10] Rice, Thomas R., "Economics of Decision Making," Research Report No. EES-DA-74-1, August 31, 1974, Dept. of Engineering-Economic Systems, Stanford University, Stanford, California.

- [11] Shuford, Emir H., "CORTEX: A Computer-Based System for Aiding Decision Making" Project 4690, Task 469003, Decision Sciences Laboratory, U.S. Air Force, L.G. Hanscom Field, Bedford. Mass., December, 1964.

- [12] Tribus, Myron "Rational Descriptions, Decisions and Designs," Pergamon Press, New York, 1969.

- [13] Zamora, Ramon M. and Ellen B. Leaf, "Tutorial on the Use of the SRI Tree Language System," Technical Memorandum, Stanford Research Institute, December 1974.